arXiv:1703.04156v1 [math.OC] 12 Mar 2017

A trust-region method for derivative-free nonlinear constrained stochastic optimization

F. Augustin[1] and Y. M. Marzouk[2]

[1]Massachusetts Institute of Technology, 77 Massachusetts Avenue, Cambridge, MA 02139. Email: fmaugust@mit.edu

[2]Massachusetts Institute of Technology, 77 Massachusetts Avenue, Cambridge, MA 02139. Email: ymarz@mit.edu

Abstract

In this work we introduce the algorithm (S)NOWPAC (Stochastic Nonlinear Optimization With Path-Augmented Constraints) for stochastic nonlinear constrained derivative-free optimization. The algorithm extends the derivative-free optimizer NOWPAC [9] to be applicable to nonlinear stochastic programming. It is based on a trust region framework, utilizing local fully linear surrogate models combined with Gaussian process surrogates to mitigate the noise in the objective function and constraint evaluations. We show several benchmark results that demonstrate (S)NOWPAC's efficiency and highlight the accuracy of the optimal solutions found.

1 Introduction

In this work we introduce a novel approach for nonlinear constrained stochastic optimization of a process, represented by an objective function f, and derive an algorithm for finding optimal process specific design parameters x within a set $X = \{x \,:\, c(x) \leq 0\} \subseteq \mathbb{R}^n$ of admissible feasible design parameters. The functions $c = (c_1, \ldots, c_r)$ are called constraint functions. Thereby we focus on methodologies that do not require gradients, only utilizing on black box evaluations of f and the constraints c. This is particularly advantageous in stochastic optimization where gradient estimation is challenging. In particular we extend the trust region optimization algorithm NOWPAC [9] by generalizing its capabilities to optimization under uncertainty with inherently noisy evaluations of the objective function and the constraints.

The main contributions of this work are threefold. Firstly, we generalize the trust region management of derivative-free optimization procedures by coupling the trust region size with the noise in the evaluations of f and c, which allows us to control the structural error in the local surrogate models used in NOWPAC. Secondly, we propose a procedure to recover feasibility of points that were falsely quantified as feasible due to the noise in the constraint evaluations. Finally, we introduce Gaussian process models of the objective function and the constraints to overall reduce the negative impact of the noise on the optimization process. We refer to Section 3 for a detailed discussion of our contributions.

Before we go into more details about our approach, we briefly introduce prototypical applications for our approach. We are concerned with optimizing objective functions f subject to constraints c that all depend on uncertain parameters $\theta \in \Theta$ not known with absolute certainty at the time of optimization; cf. [13, 53]. For example, θ may reflect limited accuracy in measurement data or it may model our lack of knowledge about process parameters. The general parameter-dependent nonlinear constrained optimization problem can be stated as

$$\begin{aligned} &\min f(x, \theta) \\ &\text{s.t.} \quad c(x, \theta) \leq 0. \end{aligned} \tag{1}$$

In general the solution of (1) depends on θ and, in order to control and limit possibly negative effects on the optimal design, we have to take the variability in θ in the optimization into account. We do this by reformulating (1) into a robust optimization problem [12, 14] using *robustness measures* $\mathcal{R}^f$ and $\mathcal{R}^c$,

$$\begin{aligned} &\min \mathcal{R}^f(x) \\ &\text{s.t.} \quad \mathcal{R}^c(x) \leq 0. \end{aligned} \tag{2}$$

We discuss a variety of robustness measures in Section 2 and refer to the rich literature on risk and deviation measures for further details; see [1, 5, 40, 58, 64, 65, 66, 78, 80]. In order to simplify notation we omit the superscripts f and c subsequently whenever the reference is clear form the context.

With only black box evaluations available, we only have access to approximations of $\mathcal{R}$ for solving (2), which therefore becomes a stochastic optimization problem. For these type of problems there exist several optimization techniques to approximate solutions. One of

them is Sample Average Approximation (SAA) [39], where a set of samples $\{\theta_i\}_{i=1}^N \subset \Theta$ is chosen for approximating the robustness measures $R_N \approx \mathcal{R}$ before starting the optimization. This set of samples is then fixed throughout the optimization process to minimize the sample approximated objective function $\mathcal{R}^f$. This procedure results in approximate solutions of (2) that depend on the particular choice of samples used. In order to reduce the associated approximation error, typically several optimization runs are averaged or the sample size $N \to \infty$ is increased; see [2, 67, 71]. An error analysis of an SAA approach for constrained optimization problems can be found in [10]. The advantage of SAA is that the inaccuracies in the sample approximation R_N of $\mathcal{R}$ are fixed to a non-random error and thus deterministic black box optimization methods can be used to solve the optimization problem.

Alternative approaches to SAA re-sample the uncertain parameters θ every time the robustness measures are evaluated. Due to the re-sampling, the evaluations of the approximate robustness measures $R_N(x)$ exhibit *sampling noise* (see Figure 1) and thus solving (2) requires stochastic optimization methods.

If the noise is small enough, i.e. the sample size N is sufficiently large, pattern search methods may be used to solve the optimization problem. Avoiding gradient approximations makes these methods less sensitive to noisy black box evaluations of the robust objective and constraints. Since the early works by Hookes and Jeeves [32] and Nelder and Mead [52, 75], there has been a significant research effort in various extensions and developments of excellent direct search optimization procedures; see for example [6, 7, 8, 24, 46, 47, 54, 55]. Alternatively, surrogate model based optimization. [9, 16, 21, 35, 48, 61, 62, 68] can be used to solve (2). Having sufficiently accurate approximations of gradients even convergence results for these methods exist; see [17, 20, 28, 42]. Here, *sufficiently accurate*, however, requires the gradient approximation to become increasingly accurate while approaching an optimal solution. This idea is incorporated in recently proposed derivative-free stochastic optimization procedures like STRONG [19] and ASTRO-DF [72], which have elaborate mechanisms to reduce the noise in black box evaluations by taking averages over an increasing number of samples while approaching an optimal design.

Thus far we only discussed optimization methods that rely on a diminishing magnitude of the noise in the robustness measure approximations and we now turn our attention to methods without this requirement. In 1951, Robbins and Monroe [63] pioneered by proposing the Stochastic Approximation (SA) method. Since then SA has been generalized to various gradient approximation schemes, e.g. by Kiefer and Wolfowitz (KWSA) [38] and Spall [73, 74, 79] with the Simultaneous Perturbation Stochastic Approximation (SPSA). We refer to [15, 36, 41] for a detailed introduction and theoretical analysis of SA methods and only remark that for all SA methods a variety of technical parameters, like the step and stencil sizes, have to be chosen very carefully. Despite a rich literature and theoretical results, this choice remains a challenging task in applying SA approaches: optimal and heuristic choices exist [74], however, they are highly problem dependent and have a strong influence on the performance and efficiency of SA methods.

Finally, Bayesian Global Optimization (BGO) [49, 50] can be used to solve (2). In BGO a global Gaussian process approximation of the objective function is build in or-

der to devise an exploration and exploitation scheme for global optimization based on expected improvement or knowledge gradients, see for example [25, 33]. We borrow the idea of building a global Gaussian surrogate for our optimization approach, however, since our focus is on local optimization we combine the Gaussian surrogate with a trust region framework. Moreover, our proposed algorithm is able to handle nonlinear constraints, a topic which only recently has has gained attention in BGO, see [27]. One particular approach, constrained Bayesian Optimization (cBO), based on expected constrained improvement optimization can be found in [26]. We discuss the advantages and disadvantages of cBO as compared to our algorithm in Section 4.2.

The outline of this paper is as follows. In Section 2 we discuss a variety of robustness measures to rigorously define the robust formulation (2). Within the same section, we also introduce sampling based approximations of robustness measures along with their confidence intervals for statistical estimation of their sampling errors. Thereafter, in Section 3, we briefly recap the trust-region algorithm NOWPAC [9] which we then generalize to make it applicable to stochastic (noisy) robust optimization tasks. We close with numerical examples in Section 4 and conclude in Section 5.

2 Robust optimization formulations

In this section we introduce a collection of robustness measures $\mathcal{R}^f$ and $\mathcal{R}^c$ to model robustness and risk within the robust optimization problem (2). Furthermore, we discuss sampling approximations of the robustness measures along with their associated confidence intervals. To simplify the notation we refer to the objective function f and the constraints c as black box b, the corresponding robustness measures will be denoted by $\mathcal{R}^b$. We assume that b is square integrable with respect to θ, i.e. its variance is finite, and that its cumulative distribution function is continuous and invertible at every fixed design point $x \in \mathbb{R}^n$.

2.1 Robustness measures

Before we start our discussion about robustness measures we point out that not all robustness measures discussed within this section are coherent risk measures in the sense of [4, 5] and we refer to [66, 76] for a detailed discussion about risk assessment strategies. However, since the focus of the present work is on optimization techniques for stochastic robust optimization, we include widely used robustness (possibly non-coherent formulations) also comprising variance terms, chance constraints and the Value at Risk. In the forthcoming we use the notation $\theta := (\theta_1, ..., \theta_m) : (\Omega, \mathcal{F}, \mathbb{P}) \to (\Theta, \mathcal{B}(\Theta), \mu)$ for the uncertain parameters mapping from a probability space $(\Omega, \mathcal{F}, \mathbb{P})$ to $(\Theta, \mathcal{B}(\Theta), \mu)$, $\Theta \subseteq \mathbb{R}^m$. Here, $B(\Theta)$ denotes the standard Borel σ-field.

Expected value and variance

The first robustness measure is the expected value

$$\mathcal{R}_0^b(x) := \mathbb{E}_\theta \left[b(x;\theta) \right] = \int_\Theta b(x;\theta) d\mu.$$

Although it may be arguable that the expected value measures *robustness* with respect to variations in θ, since it does not inform about the spread of b around $\mathcal{R}_0^b(x)$, it is a widely applied measure to handle uncertain parameters in optimization problems. For example, the expected objective value, $\mathcal{R}_0^f$, yields a design that performs best on average, whereas $\mathcal{R}_0^c$ specifies feasibility in expectation. In order to also account for the spread of realizations of b around $\mathcal{R}_0^b$ for different values of θ in a statistical sense, justifying the term *robustness* measure, a variance term,

$$\mathcal{R}_1^b(x) := \mathbb{V}_\theta \left[b(x;\theta) \right] = \mathbb{E}_\theta \left[b(x;\theta)^2 \right] - \mathcal{R}_0^b(x)^2,$$

can be included. We remark that the linear combination

$$\mathcal{R}_2^b(x) := \gamma c_1 \mathcal{R}_0^b(x) + (1-\gamma) c_2 \mathcal{R}_1^b(x),$$

$\gamma \in [0,1]$, c_1, $c_2 > 0$, of $\mathcal{R}_0^b$ and $\mathcal{R}_1^b$ has a natural interpretation in decision making. By minimizing the variance term we gain confidence in the optimal value being well represented by $\mathcal{R}_0^b$. Combining the two goals of objective minimization in expectation and the reduction of the spread of possible outcomes, the robustness measure $\mathcal{R}_2^b(x)$ provides a trade off between two possibly contradicting goals. The user's priority in one goal over the other is reflected by the weighting factor γ. The constants c_1 and c_2 are required to obtain a proper scaling between $\mathcal{R}_0^b$ and $\mathcal{R}_1^b$. Finally we remark that it is well known that $-\mathcal{R}_0^b$ is a coherent risk measure, whereas $\mathcal{R}_2^b$ is not (see [5]).

Worst case scenarios

The measure which is traditionally the most closely associated with robust optimization is the worst case formulation

$$\mathcal{R}_3^b(x) := \max_{\theta \in \Theta} \left\{ b(x,\theta) \right\}.$$

Worst case formulations are rather conservative as they yield upper bounds on the optimal objective values or, in case of worst case constraints, feasibility of the optimal design for all realizations of $\theta \in \Theta$. They are also known as hard constraints; see e.g. [12]. It is often computationally challenging to evaluate $\mathcal{R}_3^b$ and only in special cases of simple non-black box functions b it is possible to analytically compute $\mathcal{R}_3^b(x)$, which then yields a deterministic optimization problem, see f.e. [11, 29, 37, 71, 78]. In the general case of nonlinear black box functions, however, the treatment of worst case formulations requires global optimization over the uncertainty space Θ. We refer to the literature on semi-infinite programming that aims at efficiently tackle this class of problem formulations [30, 60]. Since the focus of the present work is on optimization with noisy function evaluations, we refrain from a further discussion of worst case formulations.

Probabilistic constraints and Value at Risk (VaR)

Commonly used robustness measures are probabilistic constraints, also know as chance constraints [56], where a minimal probability level $\beta \in]0,1[$ is specified up to which the optimal design has to be feasible. The corresponding measure to model allowable unlikely, with probability $1-\beta$, constraint violations is

$$\mathcal{R}_4^{b,\beta}(x) := \mu\left[b(x,\theta) \geq 0\right] - (1-\beta) = \mathbb{E}_\theta\left[\mathbb{1}(b(x,\theta) \geq 0)\right] - (1-\beta).$$

Probabilistic constraints are used to model economical aspects, e.g. to model construction costs of a power plant not exceeding a prescribed budget up to with probability β. Another example of probabilistic constraints is to constrain the gas mixture in a combustion chamber to prevent extinction of the flame with (high) probability β. A penalty for the associated costs or risks for violating the constraints can be included in the objective function. See also [43, 44, 45] for an efficient method for approximating $\mathcal{R}_4^{b,\beta}$. Under the assumption of invertible cumulative distribution functions, F_μ, of μ we have an equivalent formulation of probabilistic constraints in terms of quantile functions,

$$\mathcal{R}_5^{b,\beta}(x) := \min\left\{\alpha \in \mathbb{R} \,:\, \mu[b(x,\theta) \leq \alpha] \geq \beta\right\},$$

resulting in two equivalent formulations of the same feasible set: $\{x \in \mathbb{R}^n \,:\, \mathcal{R}_4^{c,\beta}(x) \leq 0\} = \{x \in \mathbb{R}^n \,:\, \mathcal{R}_5^{c,\beta}(x) \leq 0\}$. We will see in Section 2.2 that $\mathcal{R}_5^{b,\beta}$ often exhibits favorable smoothness properties as compared to $\mathcal{R}_4^{b,\beta}$, making it more suitable to model probabilistic constraints in our optimization procedure. We finally remark that for $b = f$, the robustness measure $\mathcal{R}_5^{f,\beta}$ is also known as Value at Risk (VaR), a widely used non-coherent risk measure in finance applications.

Conditional Value at Risk (CVaR)

The Conditional Value at Risk (CVaR) [1, 66] a coherent extension of $\mathcal{R}_5^{b,\beta}(x)$, being the conditional expectation of b exceeding the VaR:

$$\mathrm{CVaR}_\beta(x) := \frac{1}{1-\beta} \int\limits_{b(x,\theta) \geq \mathcal{R}_5^{b,\beta}(x)} b(x,\theta) d\mu.$$

Following [3, 66] we define the robustness measure

$$\mathcal{R}_6^{b,\beta}(x,\gamma) := \gamma + \frac{1}{1-\beta}\mathbb{E}_\theta\left[\left[b(x,\theta) - \gamma\right]^+\right],$$

with $[z]^+ = \max\{z, 0\}$, which allows us to minimize the CVaR without having to compute $\mathcal{R}_5^{b,\beta}$ first as minimizing $\mathcal{R}_6^b$ over the extended feasible domain $X \times \mathbb{R}$ yields

$$\min_{x \in X} \mathrm{CVaR}_\beta(x) = \min_{(x,\gamma) \in X \times \mathbb{R}} \mathcal{R}_6^b(x,\gamma).$$

2.2 Statistical estimation of robustness measures

As outlined in the previous section, the robustness measures $\mathcal{R}_0^b$, $\mathcal{R}_1^b$, $\mathcal{R}_2^b$, $\mathcal{R}_4^{b,\beta}$ and $\mathcal{R}_6^{b,\beta}$ can be written in terms of an expectation,

$$\mathcal{R}^b(x) := \mathbb{E}_\theta \left[B(x, \theta) \right], \tag{3}$$

where the function B is defined as in Table 1. Throughout this paper we assume that B

Table 1: Integrands in (3) defining the robustness measures $\mathcal{R}_0^b$, $\mathcal{R}_1^b$, $\mathcal{R}_2^b$, $\mathcal{R}_4^b$ and $\mathcal{R}_6^b$.

robustness measure	integrand $B(x, \theta)$
$\mathcal{R}_0^b(x)$	$b(x, \theta)$
$\mathcal{R}_1^b(x)$	$\left(b(x, \theta) - \mathcal{R}_0(x)\right)^2$
$\mathcal{R}_2^b(x)$	$\gamma c_1 b(x, \theta) + (1 - \gamma) c_2 \left(b(x, \theta) - \mathcal{R}_0(x)\right)^2$
$\mathcal{R}_4^{b,\beta}(x)$	$\mathbb{1}(b(x; \theta) \geq 0) - (1 - \beta)$
$\mathcal{R}_6^{b,\beta}(x, \gamma)$	$\gamma + \frac{1}{1-\beta} \left[b(x, \theta) - \gamma\right]^+$

has finite variance. Note that in case of $\mathcal{R}_0^b(x)$, $\mathcal{R}_4^{b,\beta}(x)$ and $\mathcal{R}_6^{b,\beta}(x)$ this already follows from our assumption that b is square integrable with respect to θ . However, for the variance of B in $\mathcal{R}_1^b(x)$ and $\mathcal{R}_2^b(x)$ to be finite we require the stronger integrability condition of b^2 being square integrable. For the approximation of (3) at x we use a sample average E_N based on N samples $\{\theta_i\}_{i=1}^N \sim \mu$,

$$\mathbb{E}_\theta \left[B(x, \theta) \right] = E_N \left[B(x, \theta) \right] + \varepsilon_x = \frac{1}{N} \sum_{i=1}^{N} B(x, \theta_i) + \varepsilon_x. \tag{4}$$

Here ε_x represents the error in the sample approximation. From the Central Limit Theorem we know that $\sqrt{N}\varepsilon_x$ is asymptotically normal distributed with zero mean and variance $\sigma^2 = \mathbb{V}_\theta[B(x, \theta)]$ for $N \to \infty$. This allows the definition of a confidence interval around the approximated expected value, $E_N \left[B(x, \theta) \right]$, which contains $\mathbb{E}_\theta \left[B(x, \theta) \right]$ with high probability. To get a confidence interval

$$\left[\mathbb{E}_N \left[B(x, \theta) \right] - \bar{\varepsilon}_x,\ \mathbb{E}_N \left[B(x, \theta) \right] + \bar{\varepsilon}_x \right]$$

that contains $\mathbb{E}_\theta \left[B(x, \theta) \right]$ with a probability exceeding $\nu \in]0, 1[$ we compute the sample estimate $s_N(x)$ of the standard deviation of $\{B(x, \theta_i)\}_{i=1}^N$,

$$s_N(x)^2 = \frac{1}{N-1} \sum_{i=1}^{N} \left(B(x, \theta_i) - E_N[B(x, \theta)] \right)^2, \tag{5}$$

and set

$$\bar{\varepsilon}_x = \frac{t_{N-1,\nu}\, s_N(x)}{\sqrt{N}},$$

with $t_{N-1,\nu}$ being the ν-quantile of the Student-t distribution. In our proposed Algorithm 4 we use $\bar{\varepsilon}_x$ as an indicator for the upper bound on the sampling error $\varepsilon_x \leq \bar{\varepsilon}_x$ with probability exceeding ν. We chose $t_{N-1,\nu} = 2$ in our implementation which yields a confidence level exceeding 0.975 for sample sizes $N \geq 60$.

We now come back to the discussion about the smoothness properties of the robustness measures $\mathcal{R}_4^{b,\beta}$ and $\mathcal{R}_5^{b,\beta}$. In Example 2.1 we show that $\mathcal{R}_4^{b,\beta}$ often exhibits large curvatures or even non-smoothness in x, creating a challenge for approximating this robustness measure using surrogate models. We therefore use the quantile reformulation $\mathcal{R}_5^{b,\beta}$ over the probabilistic constraints $\mathcal{R}_4^{b,\beta}$.

Example 2.1 (Non-smoothness of $\mathcal{R}_4^{b,\beta}$)**.** *Let us consider the two robust constraints*

$$\mathcal{R}_4^{c_1,\beta}(x) = \mathbb{E}_\theta\left[\mathbb{1}\left\{x \,:\, \exp\left(\frac{x}{2}\right) - 16(x-2)^2\theta^2 + x - 1 \geq 0\right\}\right] - 0.1$$
$$\mathcal{R}_4^{c_2,\beta}(x) = \mathbb{E}_\theta\left[\mathbb{1}\{x \,:\, 30x + \theta) \geq 0\}\right] - 0.1$$

with $\theta \sim \mathcal{N}(0,1)$ and $\beta = 0.9$. We compute the sample average estimator using 1000 samples and plot the robustness measures $\mathcal{R}_4^{c_1,\beta}$ (top left) and $\mathcal{R}_4^{c_2,\beta}(x)$ (bottom left) in Figure 1. Besides the sample noise we observe that the response surface of $\mathcal{R}_4^{c_1,\beta}$ has kinks at $x \approx 0$, $x \approx 1.5$ and $x \approx 2.5$ which violates the smoothness assumptions on the constraints; for an in depths discussion about smoothness properties of probability distributions we refer to [37, 77, 78]. Apart from the kinks, even in cases where $\mathcal{R}_4^{c,\beta}$ is arbitrarily smooth, cf. $\mathcal{R}_4^{c_2,\beta}$, it may be a close approximation to a discontinuous step function. The quantile formulations of the probabilistic constraints, $\mathcal{R}_5^{c_1,\beta}$ (top right) and $\mathcal{R}_5^{c_2,\beta}$ (bottom right) in Figure 1, on the other hand exhibit smooth behavior. $\diamond$

To approximate the quantile function $\mathcal{R}_5^{b,\beta}$ we can not rely on (4) for approximating $\mathcal{R}_5^{b,\beta}$ anymore. Instead we follow [81] and use the order statistic $b_{1:N}^x \leq \cdots \leq b_{N:N}^x$, $b_{i:N}^x \in \{b(x,\theta_i)\}_{i=1}^N$ to compute an approximation $b_{\bar{\beta}:N}^x$ of the quantile $b_\beta(x)$. More specifically we choose the standard estimator $b_{\bar{\beta}:N}^x \approx b_\beta(x)$ with

$$\bar{\beta} = \begin{cases} N\beta, & \text{if } N\beta \text{ is an integer and } \beta < 0.5 \\ N\beta + 1, & \text{if } N\beta \text{ is an integer and } \beta > 0.5 \\ \frac{N}{2} + \mathbb{1}(U \leq 0), & \text{if } N\beta \text{ is an integer and } \beta = 0.5 \\ \lfloor N\beta \rfloor + 1, & \text{if } N\beta \text{ is not an integer} \end{cases}$$

and $U \sim \mathcal{U}[0,1]$, yielding

$$b_\beta(x) = R_4^{b,\beta}(x) + \varepsilon_x = b_{\bar{\beta}:N}^x + \varepsilon_x. \tag{6}$$

Since the order statistic satisfies

$$\mu_b[b_{l:N}^x \leq b_\beta(x) \leq b_{u:N}^x] \geq \sum_{i=l}^{u-1} \binom{N}{i} \beta^i (1-\beta)^{N-i} =: \pi(l,u,N,\beta)$$

we use it to define a highly probable confidence interval $\left[b_{l:N}^k, b_{u:N}^k\right]$; see [23]. In the same way as for the sample averages (4) we obtain a highly probable upper bound $\bar{\varepsilon}_x$ on ε_x by choosing

$$\bar{\varepsilon}_x := \max\left\{b_{\bar{\beta}:N}^x - b_{(\bar{\beta}-i):N}^x, b_{(\bar{\beta}+i):N}^x - b_{\bar{\beta}:N}^x\right\}$$

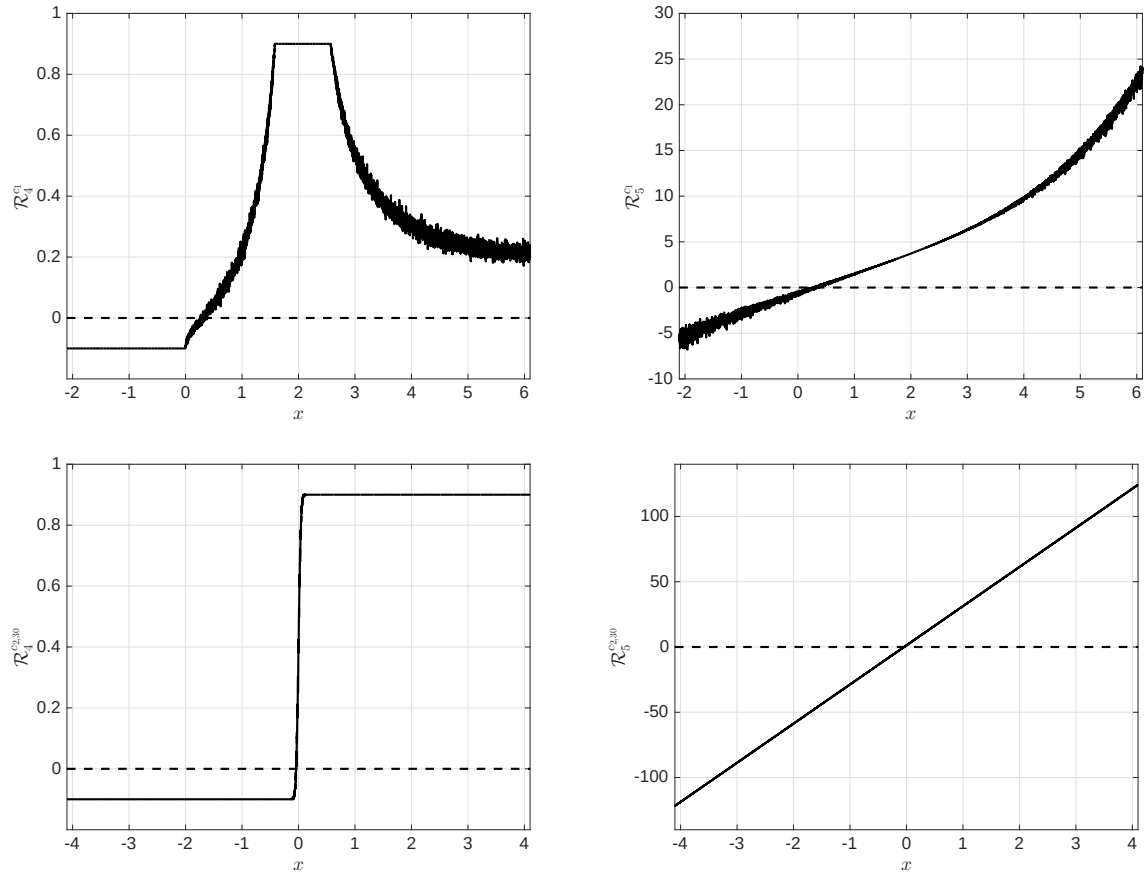

Figure 1: Sample approximation of $\mathcal{R}_4^{c_1,0.9}$ (upper left) and $\mathcal{R}_5^{c_1,0.9}$ (upper right) based on resampling 1000 samples at each x. The thresholds 0 are plotted as dashed lines. The lower plots show $\mathcal{R}_4^{c_2,0.9}$ (left) and $\mathcal{R}_5^{c_2,0.9}$ (right).

for an $i \in \{1, \dots, N\}$ such that $\pi\left(\bar{\beta} - i, \bar{\beta} + i, N, \beta\right) \geq \nu$ for the confidence level $\nu \in]0, 1[$. We refer to [82] for a detailed discussion about optimal quantile estimators.

3 Stochastic nonlinear constrained optimization

Within the last Section 2.2 we saw that all sample estimates of the robustness measures exhibit sampling noise ε_x. Therefore, we propose a stochastic optimization framework, which is based on the black box optimizer NOWPAC [9], to solve

$$\begin{aligned} &\min R_N^f(x) \\ &\text{s.t.} \quad R_N^c(x) \leq 0. \end{aligned} \tag{7}$$

Here, $R_N^b(x)$ represents any of the approximated robustness measures (4) or (6). Within the Section 3.1 we briefly review NOWPAC's key features to set the stage for its generalization to (S)NOWPAC — (Stochastic) Nonlinear Optimization With Path-Augmented Constraints — in Sections 3.2- 3.4.

3.1 Review of the trust region framework NOWPAC

NOWPAC [9] is a derivative-fee trust region optimization framework that uses black box evaluations to build fully linear (see [22]) surrogate models $m_k^{\mathcal{R}^f}$ and $m_k^{\mathcal{R}^c}$ of the objective function and the constraints within a neighborhood of the current design x_k, $k = 0, \ldots$. We define this neighborhood to be $\{x \in \mathbb{R}^n \,:\, \|x - x_k\| \leq \rho_k\}$ and call it a trust region with trust region radius $\rho > 0$. The feasible domain $X := \{x \in \mathbb{R}^n \,:\, \mathcal{R}^c(x) \leq 0\}$ is bounded by the robust constraints, where we use the short-hand notation $\mathcal{R}^c(x) := (\mathcal{R}^{c_1}(x), \ldots, \mathcal{R}^{c_r}(x))$ to denote r constraints. The optimization is performed as follows: starting from $x_0 \in X$ a sequence of intermediate points $\{x_k\}_k$ is computed by solving the trust region subproblems

$$\begin{aligned} x_k := \arg\min\, & m_k^{\mathcal{R}^f}(x) \\ \text{s.t.} \quad & x \in X_k,\ \|x - x_k\| \leq \rho_k \end{aligned} \tag{8}$$

with the approximated feasible domain

$$X_k := \left\{x \in \mathbb{R}^n \,:\, m_k^{\mathcal{R}^c}(x) + h_k(x - x_k) \leq 0\right\}. \tag{9}$$

The additive offset h_k to the constraints is called the inner boundary path, a convex offset-function to the constraints ensuring convergence of NOWPAC. We refer to [9] for more details on the inner boundary path. Having computed x_k NOWPAC only accepts this trial step if it is feasible with respect to the exact constraints $\mathcal{R}^c$, i.e. if $\mathcal{R}^c(x_k) \leq 0$. Otherwise the trust region radius is reduced and, after having ensured fully linearity of the models $m_k^{\mathcal{R}^f}$ and $m_k^{\mathcal{R}^c}$, a new trial step x_k is computed. To assess closeness to a first order optimal point we use the criticality measure

$$\alpha_k(\rho_k) := \frac{1}{\rho_k} \left| \min_{\substack{x_k + d \in X_k \\ \|d\| \leq \rho_k}} \left\langle g_k^{\mathcal{R}^f}, d \right\rangle \right|, \tag{10}$$

where $g_k^{\mathcal{R}^f} = \nabla m_k^{\mathcal{R}^f}(x_k)$ is the gradient of the surrogate model of the objective function $\mathcal{R}^f$ at x_k. We recall the simplified algorithm for NOWPAC within Algorithm 1.

3.2 Gaussian process supported trust region management

The efficiency of Algorithm 1 depends on the accuracy of the surrogate models $m_k^{\mathcal{R}^b}$ and subsequently our ability to predict a good reduction of the objective function within the subproblem (8). To ensure a good approximation quality we firstly introduce a noise-adapted trust region management to NOWPAC to couple the structural error in the surrogate approximations and the sampling error in the evaluation of R_N. Secondly we propose the construction of Gaussian processes to reduce the sampling noise in the black box evaluations to build fully-linear surrogate models of the objective function and the constraints.

Algorithm 1: Simplified NOWPAC

Construct the initial fully linear models $m_0^{\mathcal{R}^f}(x_0+s)$, $m_0^{\mathcal{R}^c}(x_0+s)$
Compute criticality measure $\alpha_0(\rho_0)$
for *k = 0, 1, ...* **do**
  `/* STEP 0: Criticality step */`
  if $\alpha_k(\rho_k)$ *is too small* **then**
    Set $\rho_k = \omega\rho_k$ and update $m_k^{\mathcal{R}^f}$ and $m_k^{\mathcal{R}^c}$
    Repeat `STEP 0`
  end
  `/* STEP 1: Step calculation */`
  Compute a trial step $s_k = \underset{x_k+s\in X_k,\, \|s\|\leq\rho_k}{\arg\min} \; m_k^{\mathcal{R}^f}(x_k+s)$
  `/* STEP 2: Check feasibility of trial point */`
  if $\mathcal{R}^c(x_k)(x_k+s_k) > 0$ **then**
    Set $\rho_k = \gamma\rho_k$ and update $m_k^{\mathcal{R}^f}$ and $m_k^{\mathcal{R}^f}$
    Go to `STEP 0`
  end
  `/* STEP 3: Acceptance of trial point and update trust region */`
  Compute $r_k = \frac{\mathcal{R}^f(x_k)-\mathcal{R}^f(x_k+s_k)}{m_k^{\mathcal{R}^f}(x_k)-m_k^{\mathcal{R}^f}(x_k+s_k)}$
  if $r_k \geq \eta_0$ **then**
    Set $x_{k+1} = x_k + s_k$
    Include x_{k+1} into the node set and update the models to $m_{k+1}^{\mathcal{R}^f}$ and $m_{k+1}^{\mathcal{R}^c}$
  else
    Set $x_{k+1} = x_k$, $m_{k+1}^{\mathcal{R}^f} = m_k^{Rf}$ and $m_{k+1}^{\mathcal{R}^c} = m_k^{\mathcal{R}^c}$
  end
  Set $\rho_{k+1} = \begin{cases} \gamma_{inc}\rho_k & \text{if} \quad r_k \geq \eta_1 \\ \rho_k & \text{if} \quad \eta_0 \leq r_k < \eta_1, \\ \gamma\rho_k & \text{if} \quad r_k < \eta_0. \end{cases}$
  Update $m_{k+1}^{\mathcal{R}^f}$ and $m_{k+1}^{\mathcal{R}^c}$
  Compute criticality measure $\alpha_k(\rho_k)$
  if $\rho_{k+1} < \rho_{min}$ **then**
    Output last design and objective value $\left(x_{k+1}, \mathcal{R}^f(x_{k+1})\right)$
    Terminate NOWPAC
  end
end

We know from, i.e., [34, Thm. 2.2] that fully linear surrogate models being constructed

using noise corrupted black box evaluations satisfy the error bound

$$\begin{aligned} \left\| \mathcal{R}^b(x_k + s) - m_k^{R^b}(x_k + s) \right\| &\leq \kappa_1 \, \rho_k^2 \\ \left\| \nabla\mathcal{R}^b(x_k + s) - \nabla m_k^{R^b}(x_k + s) \right\| &\leq \kappa_2 \, \rho_k, \end{aligned} \tag{11}$$

with high probability ν for constants

$$\kappa_i = \kappa_i \left(\bar{\varepsilon}_{max}^k \rho_k^{-2} \right), \qquad i \in \{1, 2\}. \tag{12}$$

The constants κ_1 and κ_2 depend on the poisedness constant $\Lambda \geq 1$ as well as on the estimates of the statistical upper bounds for the noise term, $\bar{\varepsilon}_{max}^k = \max_{i=1}^{\bar{n}} \{\bar{\varepsilon}_i\}$ from Section 2.2. In the presence of noise, i.e. $\bar{\varepsilon}_{max}^k > 0$, the term $\bar{\varepsilon}_{max}^k \rho_k^{-2}$ and thus κ_1 and κ_2 grow unboundedly for a shrinking trust region radius, violating the fully linearity property of $m_k^{R^b}$. Thus, in order to ensure fully linearity of the surrogate models, we have to enforce an upper bound on the error term $\bar{\varepsilon}_{max}^k \rho_k^{-2}$. We do this by imposing the lower bound

$$\bar{\varepsilon}_{max}^k \rho_k^{-2} \leq \lambda_t^{-2}, \quad \text{resp.} \qquad \rho_k \geq \lambda_t \sqrt{\bar{\varepsilon}_{max}^k}, \tag{13}$$

on the trust region radii for a $\lambda_t \in]0, \infty[$. We adapt the trust region management in NOWPAC to Algorithm 2 to guarantee (13). This noise-adapted trust region management

Algorithm 2: Noise adapted updating procedure for trust region radius.

```
Input: Trust region factor s ∈ {γ, γ_inc, θ}.
Set ρ_{k+1} = max{ s ρ_k, λ_t sqrt(ε̄^k_max) }
if ρ_{k+1} > ρ_max then
    Set ρ_{k+1} = ρ_max
end
```

couples the structural error of the fully linear approximation with the highly probable upper bound on the error in the approximation of the robustness measures. This coupling, however, also prevents the trust region radii from converging to 0, therefore limiting the level of accuracy of the surrogate models $m_k^{\mathcal{R}^b}$ and thus the accuracy of the optimization result.

In order to increase the accuracy of the optimization result, we therefore have to reduce the magnitude of the noise term $\bar{\varepsilon}_{max}^k$. To do this we introduce Gaussian process (GP) surrogates of $\mathcal{R}^b$ by using the information at points $\{(x_k^{(i)}, R_i^b)\}_{i=1}^{\hat{n}}$ where we already have evaluated the objective function and the constraints. We denote the corresponding GP surrogate models with $\mathcal{G}_b^k(x)$ and use stationary square-exponential kernels

$$K_b(x, y) = \sigma_b^2 \prod_{i=1}^{n} e^{-\frac{1}{2}\left(\frac{\|x-y\|}{l_i^b}\right)^2} \tag{14}$$

with standard deviations σ_b and length scales $l_1^b, \ldots, l_n^b$, where b again indicates either the objective function or one of the constraints. For the construction of the GPs we only take

points with a distance smaller than $2\rho_k$ to the current design point x_k. This focuses our approximation to a localized neighborhood and the stationarity assumption does not pose a severe limitation to the quality of the GP approximations. We remark, however, that other kernels may be employed to account for potentially available additional information such as non-stationarity of the objective or constraint functions. Having this Gaussian approximation we gradually replace the noisy function evaluations R_i^b with the mean of the Gaussian process surrogate,

$$\begin{aligned} \hat{R}_{k,i}^b &= \gamma_s^i \mathcal{G}_b^k(x_k^{(i)}) + (1-\gamma_s^i) R_i^b \\ \hat{\varepsilon}_{k,i}^b &= \gamma_s^i t_{N-1,\beta} \sigma_b(x_k^{(i)}) + (1-\gamma_s^i)\bar{\varepsilon}_i^b \end{aligned} \tag{15}$$

where $\sigma_b(x_k^{(i)})$ denotes the standard deviation of $\mathcal{G}_b^k$ at point $x_k^{(i)}$. The weight factor

$$\gamma_s^i := e^{-\sigma_b\left(x_k^{(i)}\right)}$$

is chosen to approach 1 when the Gaussian process becomes more and more accurate. The corrected evaluations $\hat{R}_{k,i}^b$ as well as the associated noise level $\hat{\varepsilon}_i^b$ at the local interpolation points are then used to build the local surrogate models $m_k^{R^f}$ and $m_k^{R^c}$. The intention behind using a Gaussian process surrogate model is to exploit posterior consistency of the GP for $\sigma_b(x_k^{(i)})$ converging to zero for an increasing number of evaluations of R^b within a neighborhood of x_k. In the limit the Gaussian process mean converges to the exact function $\mathcal{R}^b$ and the lower bound (13) on the trust region radius vanishes, allowing for increasingly accurate optimization results.

By combining the two surrogate models we balance two sources of approximation errors. On the one hand, there is the structural error in the approximation of the local surrogate models, cf. (11) which is controlled by the size of the trust region radius. On the other hand, we have the inaccuracy in the GP surrogate itself which is reflected by the standard deviation of the GP. Note that Algorithm 2 relates these two sources of errors by coupling the size of the trust region radii to the size of the credible interval through (15), only allowing the trust region radius to decrease if $\sigma_b(x_k^{(i)})$ becomes small.

We ensure posterior consistency[1], and thus $\sigma_b(x_k^{(i)})$ becoming smaller as x_k approaches the optimal design, in two ways: at first we observe that the increasing number of black box evaluations performed by the optimizer during the optimization process helps to increase the quality of the Gaussian process approximation. However, these evaluations may be localized and geometrically not well distributed around the current iterate x_k. We therefore enrich the set black box evaluations by randomly sampling a point

$$\tilde{x} \sim \mathcal{N}\left(x_k, \frac{3\sqrt{\rho_k}}{10} I\right),$$

to improve the geometrical distribution of the regression points for the GP surrogates whenever a trial point is rejected. This can happen either when it appears to be infeasible under the current Gaussian process approximation-corrected black box evaluations (15), or whenever if gets rejected in `STEP 3` in Algorithm 1. Secondly, we progressively re-estimate

[1] to be checked: is that correct?

the hyper-parameters in the Gaussian process regression to avoid problems with overfitting [59, 18]. To compute the correlation lengths, $\{l_i^b\}_{i=1}^n$, and standard deviations, σ_b, we maximize the marginal likelihood [59]. In the present implementation of (S)NOWPAC this parameter estimation can be triggered either at user-prescribed numbers of black-box evaluations or after $\lambda_k \cdot n$ consecutive rejected or infeasible trial steps, where λ_k is a user prescribed constant.

3.3 Relaxed feasibility requirement

An integral part of Algorithm 1 is the feasibility requirement in `STEP 2`. It checks and guarantees feasibility of all intermediate design points x_k. Checking feasibility in the presence of noise in the evaluations of constraints, however, is challenging. Thus, it might happen that Algorithm 3 - part 1 accept points, which only later, after the GP correction (15), may be revealed as being infeasible. To generalize NOWPAC's capabilities to recover from infeasible points we introduce a feasibility restoration mode. We thus have the two operational modes

(M1) objective minimization and
(M2) feasibility restoration.

The optimizer operates in mode (M1) whenever the current point x_k appears to be feasible under the current evaluations (15), whereas it switches to mode (M2) if x_k becomes infeasible and vice versa. For the definition of the modes (M1) and (M2) we simply exchange the underlying trust region subproblem to be solved: in mode (M1) the standard subproblem

$$\begin{aligned} & \min m_k^{R^f}(x_k + s_k), \\ \text{s.t.} \quad & \bar{m}_k^{R^{c_i}}(x_k + s_k) \leq 0, \ i = 1 \dots r \\ & \|s_k\| \leq \rho_k \end{aligned} \tag{16}$$

is solved for the computation of a new trial point $x_k + s_k$, where $\bar{m}_c^{R^{c_i}}$ denote the inner-boundary path augmented models of the noisy evaluations of $\mathcal{R}^{c_i}$; cf. (9). The subproblem

$$\begin{aligned} & \min \left\langle g_k^{R^f}, s_k \right\rangle, \\ \text{s.t.} \quad & \bar{m}_k^{R^{c_i}}(x_k + s_k) \leq 0, \ i = 1 \dots r \\ & \|s_k\| \leq \rho_k \end{aligned} \tag{17}$$

is used for computation of the criticality measure α_k.

In mode (M2) the subproblem

$$\begin{aligned} & \min \sum_{i \in \mathcal{I}_k} \left(m_k^{R^{c_i}}(x_k + s_k)^2 + \lambda_g m_k^{R^{c_i}}(x_k + s_k) \right), \\ \text{s.t.} \quad & \bar{m}_k^{R^{c_i}}(x_k + s_k) \leq \tau_i, \ i = 1 \dots r \\ & \|s_k\| \leq \rho_k \end{aligned} \tag{18}$$

is solved for the computation of a new trial point $x_k + s_k$, along with

$$\begin{aligned} \min \sum_{i \in \mathcal{I}_k} \left(2m_k^{R^{c_i}}(x_k) + \lambda_g\right) \left\langle g_k^{m_k^{R^{c_i}}}, s_k \right\rangle, \\ \text{s.t.} \quad \bar{m}_k^{R^{c_i}}(x_k + s_k) \leq \tau_i, \ i = 1 \dots r \\ \|s_k\| \leq \rho_k \end{aligned} \tag{19}$$

for computation of the corresponding criticality measure. Here,

$$\mathcal{I}_k = \{i \ : \ \hat{R}_k^{c_i} > 0, \ i = 1, \dots, r\},$$

with the Gaussian process corrected constraint evaluations (15) at the current point x_{k-1} The slack variables $\tau := (\tau_1, \dots, \tau_r)$ are set to $\tau_i = \max\{\hat{R}_k^{c_i}, \ 0\}$. We introduce the parameter $\lambda_g \geq 0$ in (18) and (19) to guide the feasibility restoration towards the interior of the feasible domain.

3.4 The stochastic trust region algorithm (S)NOWPAC

We now state the final Algorithm 4 of (S)NOWPAC. The general procedure follows closely the steps in Algorithm 1 and includes the generalizations we introduced in Sections 3.2 and 3.3 to handle noisy black box evaluations. A summary of all default values for internal parameters we used in our implementation of (S)NOWPAC is given in Table 2.

Table 2: Internal parameters of (S)NOWPAC and their default values

description	parameter	default value
factor for lower bound on trust region radii	λ_t	$\sqrt{2}$
poisedness threshold	Λ	100
gradient contribution to feasibility restoration	λ_g	10^{-4}

4 Numerical examples

We first discuss a two dimensional test problem in Section 4.1 to visualize the optimization process and the effect of the Gaussian process to reduce the noise. Thereafter, in Section 4.2 we discuss numerical results for (S)NOWPAC on nonlinear optimization problems from the CUTEst benchmark suite, in particular, benchmark examples from [31, 69, 70]. We use three different formulations with various combinations of robustness measures from Section 2 and the data profiles proposed in [51] to compare (S)NOWPAC with cBO, COBYLA, NOMAD as well as the stochastic approximation methods SPSA and KWSA. Since COYBLA and NOMAD are not designed for stochastic optimization they will perform better for smaller noise levels. We therefore vary the sample sizes to discuss their performance based on different magnitudes of the noise in the sample approximations of the robust objective function and constraints.

Algorithm 3 - part 1: (S)NOWPAC

Construct the initial fully linear models $m_0^{R^f}(x_0+s)$, $m_0^{R^c}(x_0+s)$
Set $x_{best} = x_0$ and $R^f_{best} = R^f_N(x_0)$
for *k = 0, 1, ...* **do**
    `/* STEP 0: Criticality step */`
    if $\alpha_k(\rho_k)$ *is too small* **then**
        Call Algorithm 2 with $\kappa = \omega$.
        Evaluate $R^f_N(x)$ and $R^c_N(x)$ for a randomly sampled x
        in a neighborhood of $B(x_k, \rho_k)$.
        Update Gaussian processes and black box evaluations.
        Construct surrogate models $m_k^{R^f}(x_k+s)$, $m_k^{R^c}(x_k+s)$.
        if *current point,* x_k*, is infeasible* **then**
            Switch to mode (M2).
        else
            Switch to mode (M1).
        end
        Repeat `STEP 0`.
    end
    `/* STEP 1: Step calculation */`
    Compute a trial step according to (16) or (18).
    Evaluate $R^f_N(x_k+s_k)$ and $R^c_N(x_k+s_k)$.
    Update Gaussian processes and black box evaluations.
    `/* STEP 2: Check feasibility of trial point */`
    if $R^{c_i}_N(x_k+s_k) > \tau_i$ *for an* $i = 1, \ldots, r$ **then**
        Call Algorithm 2 with $\kappa = \theta$.
        if *trial step* $x_k + s_k$ *is infeasible* **then**
            Evaluate $R^f_N(x)$ and $R^c_N(x)$ for a randomly sampled x
            in a neighborhood of $B(x_k, \rho_k)$.
            Update Gaussian processes and black box evaluations.
            Construct surrogate models $m_k^{R^f}(x_k+s)$, $m_k^{R^c}(x_k+s)$.
        else
            Go to `STEP 1`.
        end
        if $R^c_N(x_k) < 0$ **then**
            Switch to mode (M1).
        else
            Switch to mode (M2).
        end
    end
end

Algorithm 3 - part 2: (S)NOWPAC

```
   /* STEP 3: Acceptance of trial point and update trust region      */
 1 if acceptance ratio, r_k, is greater than η_0 then
 2 |   Set x_{k+1} = x_k + s_k
 3 |   Include x_{k+1} into the node set and update the models to m_{k+1}^{R^f} and m_{k+1}^{R^c}
 4 |   Call Algorithm 2 with κ ∈ {1, γ_inc}
 5 else
 6 |   Set x_{k+1} = x_k, m_{k+1}^{R^f} = m_k^{R^f} and m_{k+1}^{R^c} = m_k^{R^c}
 7 |   Call Algorithm 2 with κ = γ
 8 |   Evaluate R_N^f(x) and R_N^c(x) for a randomly sampled x
 9 |   in a neighborhood of B(x_k, ρ_k).
10 |   Update Gaussian processes and black box evaluations.
11 |   Construct surrogate models m_k^{R^f}(x_k + s), m_k^{R^c}(x_k + s).
12 end
13 if k = n_max then
14 |   Stop.
15 end
```

4.1 A two dimensional test example

We consider the optimization problem

$$\begin{aligned} \min \mathbb{E}\left[\sin(x-1+\theta_1)+\sin\left(\frac{1}{2}y-1+\theta_1\right)^2\right]+\frac{1}{2}\left(x+\frac{1}{2}\right)^2-y \\ \text{s.t.} \quad \mathbb{E}\left[-4x^2(1+\theta_2)-10\theta_3\right] \leq 25-10y \\ \mathbb{E}\left[-2y^2(1+\theta_4)-10(\theta_4+\theta_2)\right] \leq 20x-15 \end{aligned} \tag{20}$$

with $\theta = (\theta_1, \ldots, \theta_4) \sim \mathcal{U}[-1,1]^4$ and the starting point $x_0 = (4,3)$. For the approximation of the expected values we use $N = 50$ samples of θ and we estimate the magnitudes of the noise terms as described in Section 2.2. The noise in the objective function and constraints can be seen Figure 2. The feasible domain is to the right of the exact constraints which are indicated by dotted red lines. We see that the noise is the largest in the region around the optimal solution (red cross).

To show the effect of the noise reduction we introduced in Section 3.2, we plot the objective function and the constraints corrected by the respective Gaussian process surrogates around the current design point at 20 (upper left), 40 (upper right) and 100 (lower plots) evaluations of the robustness measures. We see that the noise is reduced which enables (S)NOWPAC to efficiently approximate the optimal solution.

Note that the approximated GP-corrected feasible domains within the trust region show significantly less noise than outside of the trust region. Moreover, we see that the optimizer eventually gathers more and more black box evaluations, yielding an increasingly better noise reduction. Looking at the noisy constraint contours at 40 evaluations, we see that the quantification of feasibility based on the Gaussian process supported black box

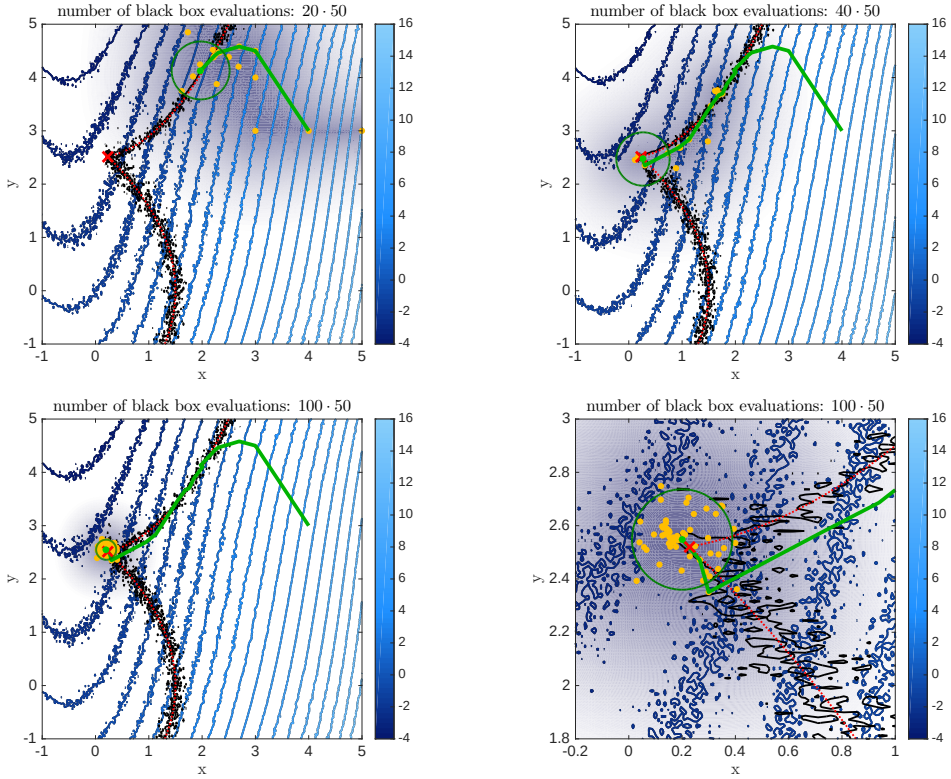

Figure 2: Realizations of the contour plots of the noisy objective function and constraints for optimization problem (20). The exact constraints are indicated by a red dotted line and the exact optimal point is marked with a red cross. The plots show the best point (green dot) and the optimization path (green line) after 20, 40 and 100 evaluations of the robustness measures; the lower right plot is zoomed in to the neighborhood of the optimal point. The corresponding trust regions are indicated by the green circle. Within the trust region we show the local smoothing effect of the Gaussian process corrected objective function and constraints. The gray cloud indicates the size weighing factor γ_s^i from (15); the darker the area the more weight is given to the Gaussian process mean. The Gaussian regression points are indicated by yellow dots.

evaluations is not always reliable. This underlines the necessity of the feasibility restoration mode we introduced in Section 3.3, which allows the optimizer to recover feasibility from points that appear infeasible.

4.2 Optimization performance on benchmark test set

Its utilization of Gaussian process surrogate models relates (S)NOWPAC to the successful class of Bayesian optimization techniques [49, 50], and its extensions for nonlinear optimization using either an augmented Lagrangian approach [27] or expected constrained improvement in the constrained Bayesian optimization (cBO) [26]. As opposed to Bayesian optimization, (S)NOWPAC introduces Gaussian process surrogates to smooth local trust region steps instead of aiming at global optimization. We will demonstrate that the combination of fast local optimization with a second layer of *smoothing* Gaussian process models makes (S)NOWPAC an efficient and accurate optimization technique. Addition-

ally, we compare the performance of (S)NOWPAC to the optimization codes COBYLA and NOMAD as well as to the stochastic approximation methods SPSA and KWSA.

We test the performances of all optimizers on the Schittkowski optimization benchmark set [31, 70], which is part of the CUTEst benchmark suit for nonlinear constraint optimization. The dimensions of the feasible domains within our test set range from 2 to 16 with a number of constraints ranging from 1 to 10. Since the problems are deterministic, we add noise to the objective functions, $f(x) + \theta_1$ and constraints, $c(x) + \theta_2$ with $(\theta_1, \theta_2) \sim \mathcal{U}[-1, 1]^{1+r}$ and solve the following three classes of robust optimization problems:

1. Minimization of the average objective function subject to the constraints being satisfied in expectation:

$$\begin{aligned} &\min \mathcal{R}_0^f(x) \\ &\text{s.t.} \quad \mathcal{R}_0^c(x) \leq 0. \end{aligned} \tag{21}$$

2. Minimization of the average objective function subject to the constraints being satisfied in 95% of all cases:

$$\begin{aligned} &\min \mathcal{R}_0^f(x) \\ &\text{s.t.} \quad \mathcal{R}_5^{c,0.95}(x) \leq 0. \end{aligned} \tag{22}$$

3. Minimization of the 95%-CVaR of the objective function subject to the constraints being satisfied on average:

$$\begin{aligned} &\min \mathcal{R}_6^{f,0.95}(x) \\ &\text{s.t.} \quad \mathcal{R}_0^c(x) \leq 0, \end{aligned} \tag{23}$$

For the approximation of the robustness measures we use three different sample sizes $N \in \{200, 1000, 2000\}$ to show the effect of reduced noise on the optimization results. We illustrate the noise magnitudes for the different sample sizes exemplarily in Figure 3. The rows show histograms of noise realizations for the sample approximations of the robustness measures $\mathcal{R}_0^b$, $\mathcal{R}_5^{b,0.95}$ and $\mathcal{R}_6^{b,0.95}$ for $b(x, \theta) = \theta$ and $\theta \sim \mathcal{U}[-1, 1]$ respectively for sample sizes $N = 200$ (left column), $N = 1000$ (middle column) and $N = 2000$ (right column). We see that the confidence regions shrink with increasing sample size for all three robustness measures. However, the reduction of the noise magnitudes becomes increasingly slower; for the expected values it is only $N^{-\frac{1}{2}}$. We will see below that one benefit of the GP based noise reduction (15) is a more rapid noise reduction by borrowing information from neighboring design points.

For the performance comparison we use a total number of optimization runs $3 \cdot 8 \cdot 100 = 2400$ (3 robust formulations, 8 benchmark problems with 100 repeated optimization runs)) and denote the benchmark set by $\mathcal{P}$. To obtain the data profiles we determine the minimal number $t_{p,S}$ of optimization steps a solver S requires to solve problem $p \in \mathcal{P}$ under the accuracy requirement

$$\frac{\left|\mathcal{R}^f(x_k) - \mathcal{R}^f(x^*)\right|}{\max\{1, |\mathcal{R}^f(x^*)|\}} \leq \varepsilon_f \quad \text{and} \quad \max_{i=1}^{r} \left\{[\mathcal{R}^{c_i}(x_k)]^+\right\} \leq \varepsilon_c.$$

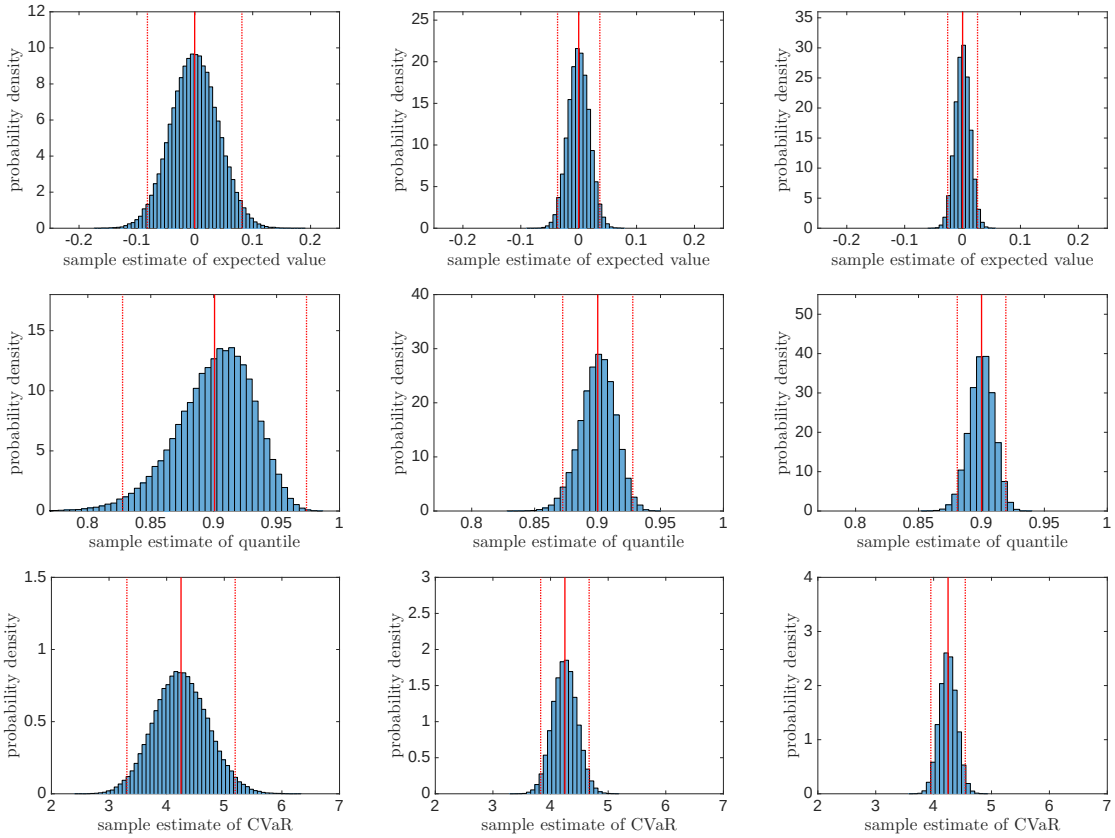

Figure 3: Histograms of $100,000$ sample estimators for the mean (upper plots), quantile (middle plots) and CVaR (lower plots) of a uniformly distribution $\mathcal{U}[-1,1]$ distributed random variable based on 200 (left plots) 1000 (middle plots) and 2000 (right plots) samples. The red line represents the mean and the dotted lines indicate the 95% confidence region. For the CVaR we show the sample approximation of the robustness measure $\mathcal{R}_6^b(x, 0.5)$, with $b(x,\theta) = \theta \sim \mathcal{U}[-1,1]$.

Hereby we limit the maximal number of optimization steps to 250 and set $t_{p,S} = \infty$ if the accuracy requirement is not met after $250 \cdot N$ black box evaluations. To decide whether the accuracy requirement is met, we use the *exact* objective and constraint values of the robustness measures which we obtained in a post-processing step. Specifically, we use the data profile

$$d_S(\alpha) = \frac{1}{2400} \left| \left\{ p \in \mathcal{P} \ : \ \frac{t_{p,S}}{n_p + 1} \leq \alpha \right\} \right|,$$

where n_p denotes the number of design parameters in problem p. We remark that, although this allows us to eliminate the influence of the noise on the performance evaluation, it is information that is not available in general. For this reason, we also include a more detailed analysis of individual optimization results below. Figure 4 shows the data profiles for different error thresholds $\epsilon_f \in \{10^{-2}, 10^{-3}\}$ and $\epsilon_c \in \{10^{-2}, 10^{-3}\}$ and for (S)NOWPAC (green), cBO (pink), COBYLA (purple), NOMAD (blue), SPSA (orange) and KWSA (red) respectively. We see that (S)NOWPAC shows a better performance than all other optimizers considered in this comparison, meaning (S)NOWPAC solves the most test

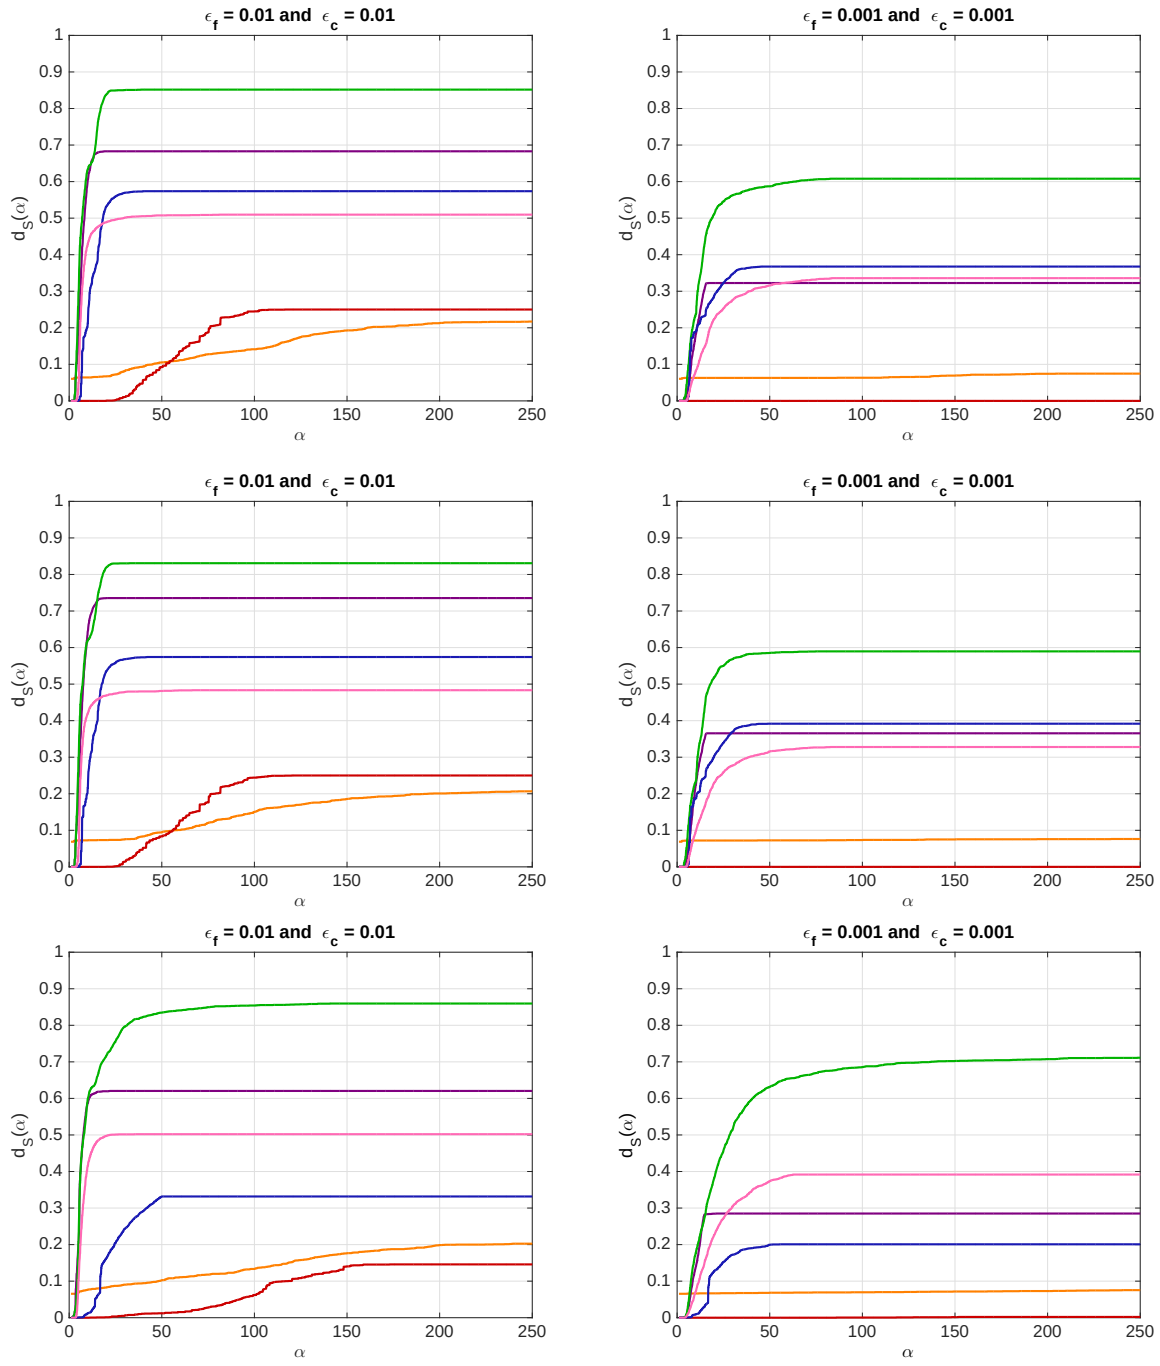

Figure 4: Data profiles for (S)NOWPAC (green), cBO (pink), COBYLA (purple), NOMAD (blue), SPSA (orange) and KWSA (red) of 2400 runs of the benchmark problems. The results for (21), (22) and (23) are plotted in the first, second and third row respectively. The profiles shown are based on the exact values for the objective function and constraints evaluated at the intermediate points computed by the respective optimizers. The data profiles are shown for varying thresholds $\epsilon_f \in \{10^{-2}, 10^{-3}\}$ and $\epsilon_c \in \{10^{-2}, 10^{-3}\}$ on the objective values and the constraint violation respectively.

problems within the given budget of black box evaluations. Looking at the performance for small values of α we also see that (S)NOWPAC exhibits a comparable or superior performance, indicating fast initial descent which is highly desirable in particular if the evaluations of the robustness measures is computationally expensive. The performance of cBO suffers in the higher dimensional benchmark problems. Here in particular the global optimization strategy of cBO naturally requires more function evaluations. Furthermore, we used the stationary kernel (14), which may not properly reflect the properties of the objective functions and constraints. A problem dependent choice of kernel function might help to reduce this problem, however, this information is often hard or even impossible to obtain in black box optimization. With the localized usage of Gaussian process approximations, see Section 3.2, (S)NOWPAC reduces the problem of violated stationarity assumptions on the objective function and constraints. As expected, COBYLA and NOMAD perform well for larger thresholds that are of the same magnitudes as the noise term in some test problems. The noise reduction in (S)NOWPAC using the Gaussian process support helps to approximate the optimal solution more accurately, resulting in better performance results. The Stochastic Approximation approaches SPSA and KWSA, despite a careful choice of hyper-parameters, do not perform well on the benchmark problems. This may be explained by the limited number of overall optimization iterations not being sufficient to achieve a good approximation of the optimal solution using inaccurate gradients.

We now show a detailed accuracy comparison of the individual optimization results at termination at $250 \cdot N$ black box evaluations. This is in contrast to the optimization results we used to compute the data profiles in Fig. 4 and reflects that in general we cannot extract the *best* design point in a post-processing step. Note that small objective values may result from infeasible points being falsely quantified as feasible due to the noise in the constraint evaluations. In Fig. 5 - 10 we therefore show the qualitative accuracy of the optimization results at the approximated optimal points at termination of the optimizers. The plot show the errors in the objective values, the constraint violations and the errors in the approximated optimal designs proposed by the optimizers at termination respectively. Since the optimal solution for test problem 268 is zero, we show the absolute error for this test problem. We use MATLAB's box plots to summarize the results for 100 optimization runs for each benchmark problem for different sample sizes $N \in \{200, 1000, 2000\}$ separately for each individual robust formulation (21)-(23). The exact evaluation of the robust objective function and constraints at the approximated optimal designs are shown to again eliminate the randomness in the qualitative accuracy of the optimization results. We see that (S)NOWPAC most reliably finds accurate approximations to the exact optimal solutions. Note that all optimizers benefit from increasing the number of samples for the approximation of the robustness measures. In (S)NOWPAC, however, the Gaussian process surrogates additionally exploit information from neighboring points to further reduce the noise, allowing for a better accuracy in the optimization results. We see that the designs computed by (S)NOWPAC and cBO match well for low-dimensional problems 29, 227, 228, but the accuracy of the results computed by cBO begins to deteriorate in dimensions larger than 4. This has two reasons: firstly, the global search strategy aims at variance reduction within the whole search domain. This requires more function evaluations than local search. Secondly, the global nature of the Gaussian processes requires a

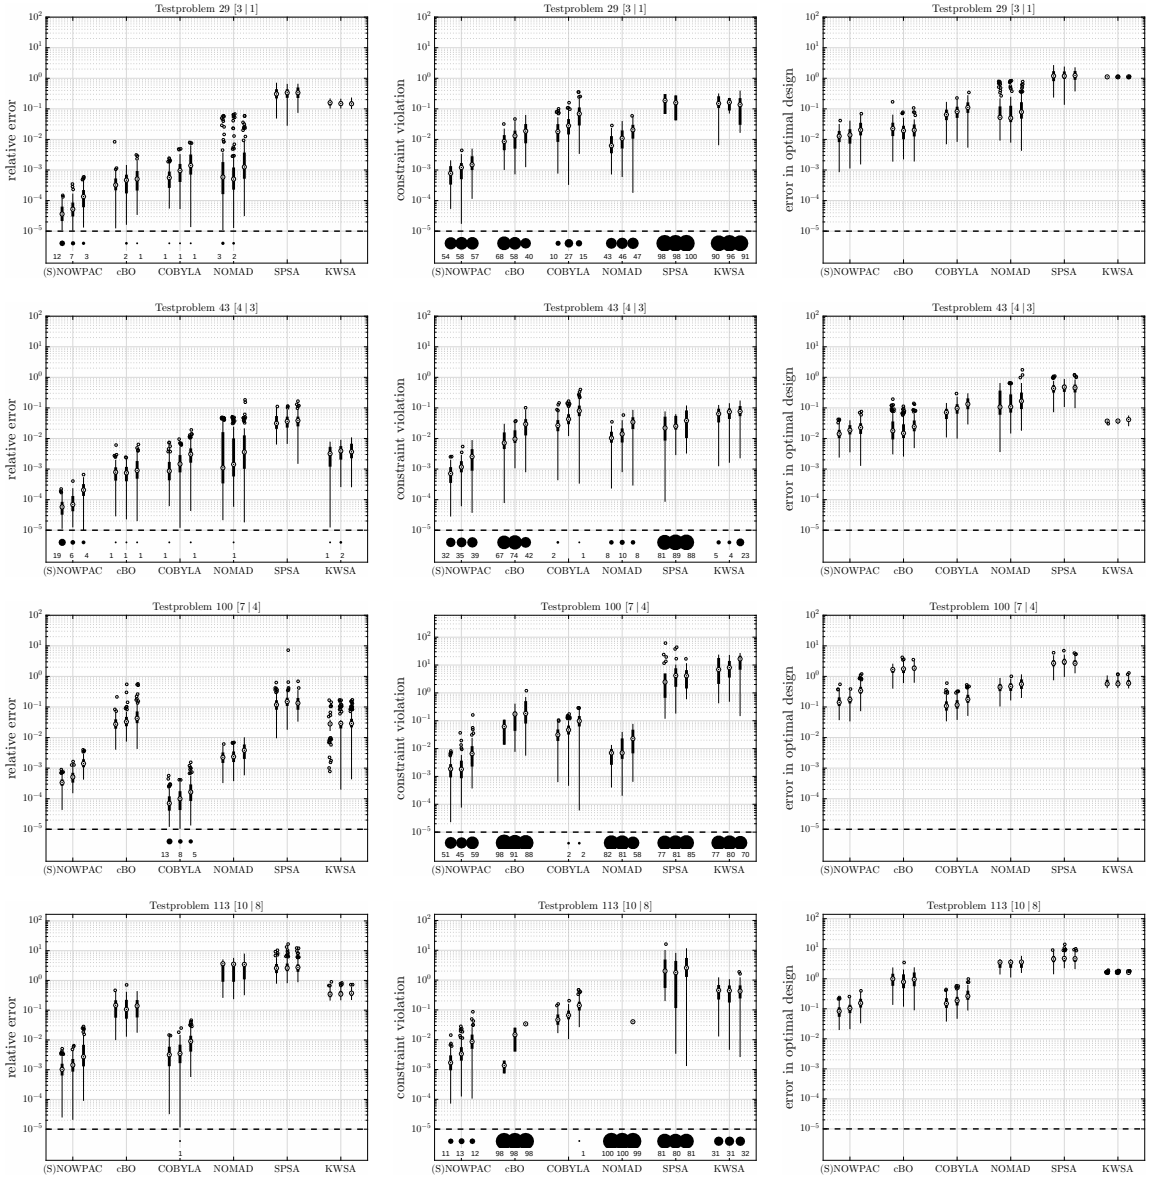

Figure 5: Box plots of the errors in the approximated optimal objective values (left plots), the constraint violations (middle plots) and the l_2 distance to the exact optimal solution (right plots) of 100 repeated optimization runs for the Schittkowski test problems number 29, 43, 100, and 113 for (21). The plots show results of the exact objective function and constraints evaluated at the approximated optimal design computed by (S)NOWPAC, cBO, COBYLA, NOMAD, SPSA and KWSA. Thereby all errors or constraint violations below 10^{-5} are stated separately below the 10^{-5} threshold and the box plots only contain data above this threshold.

suitable choice of kernels that fits to the properties of the optimization problems, i.e. non-stationarity of the optimization problem, which is not the case in all benchmark problems. Additionally, global maximization of the expected constrained improvement function in every step of the optimization procedure becomes very costly and becomes significant for more than 250 design points where the Gaussian process evaluation becomes a dominant

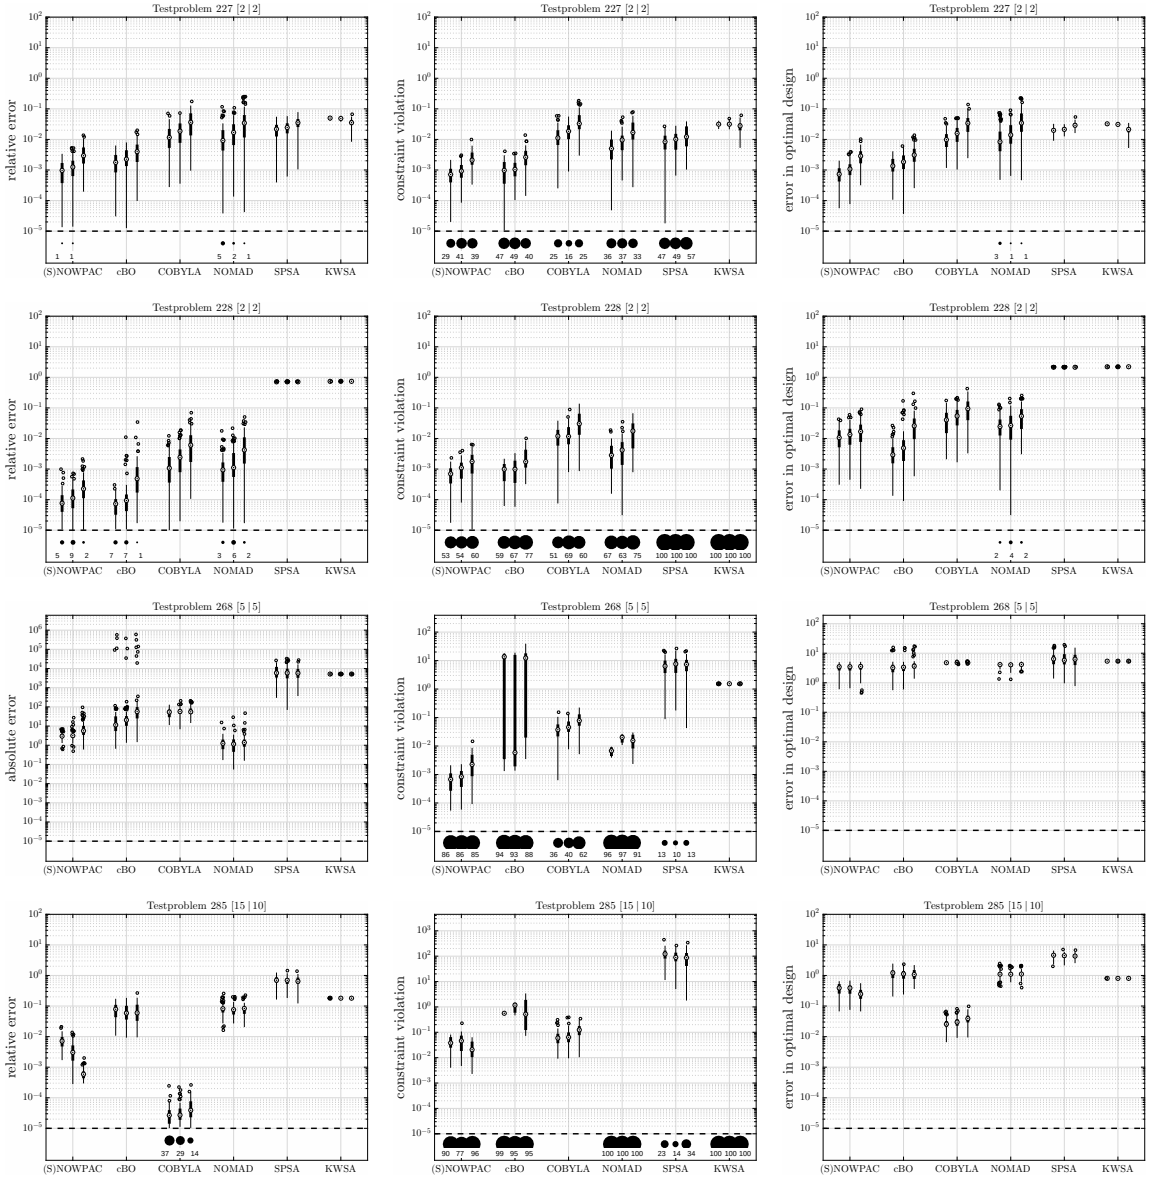

Figure 6: Box plots of the errors in the approximated optimal objective values (left plots), the constraint violations (middle plots) and the l_2 distance to the exact optimal solution (right plots) of 100 repeated optimization runs for the Schittkowski test problems number 227, 228, 268, and 285 for (21). The plots show results of the exact objective function and constraints evaluated at the approximated optimal design computed by (S)NOWPAC, cBO, COBYLA, NOMAD, SPSA and KWSA. Thereby all errors or constraint violations below 10^{-5} are stated separately below the 10^{-5} threshold and the box plots only contain data above this threshold.

source of computational effort. To reduce computational costs, approximate Gaussian processes [57] could be employed, an improvement that both, cBO and (S)NOWPAC, would benefit from. We see that, despite our careful tuning the hyper-parameters for the SPSA and KWSA approaches, the results of these optimizers are not satisfactory in most test examples. The middle plots in Fig. 5 - 10 show the maximal constraint violations

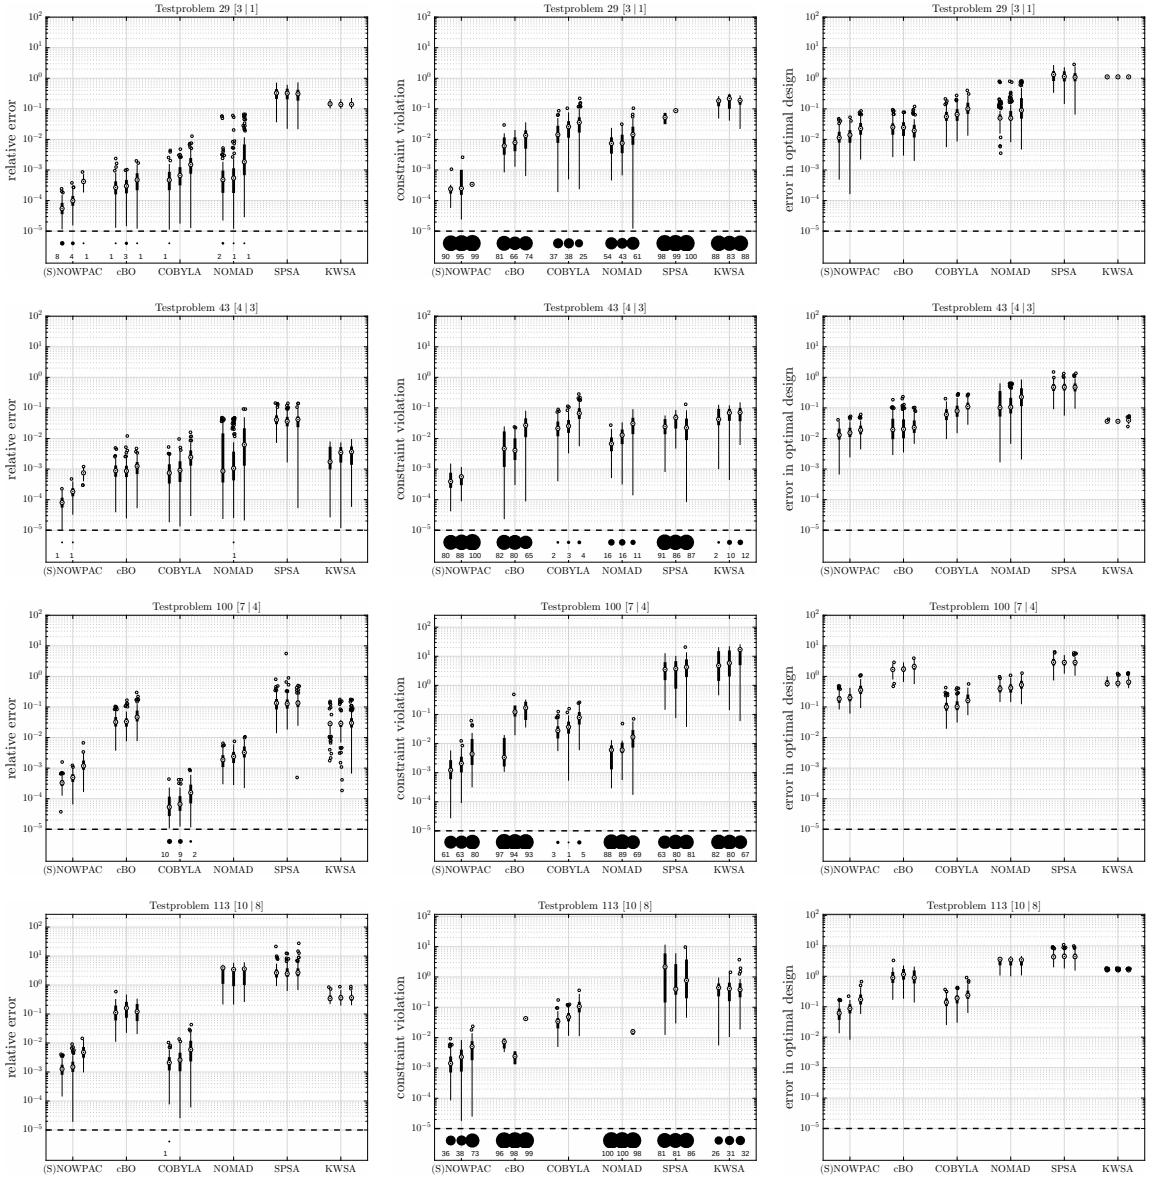

Figure 7: Box plots of the errors in the approximated optimal objective values (left plots), the constraint violations (middle plots) and the l_2 distance to the exact optimal solution (right plots) of 100 repeated optimization runs for the Schittkowski test problems number 29, 43, 100, and 113 for (22). The plots show results of the exact objective function and constraints evaluated at the approximated optimal design computed by (S)NOWPAC, cBO, COBYLA, NOMAD, SPSA and KWSA. Thereby all errors or constraint violations below 10^{-5} are stated separately below the 10^{-5} threshold and the box plots only contain data above this threshold.

at the approximated optimal designs. Here, (S)NOWPAC's constraint handling, see [9], in combination with the feasibility restoration mode from Section 3.3 allows the computation of approximate optimal designs that exhibit only small constraint violations well below the noise level; cf. Fig. 3. Finally, the right plots in Figures 5 - 10 show the error in the approximated optimal designs. We see that (S)NOWPAC yields either comparable

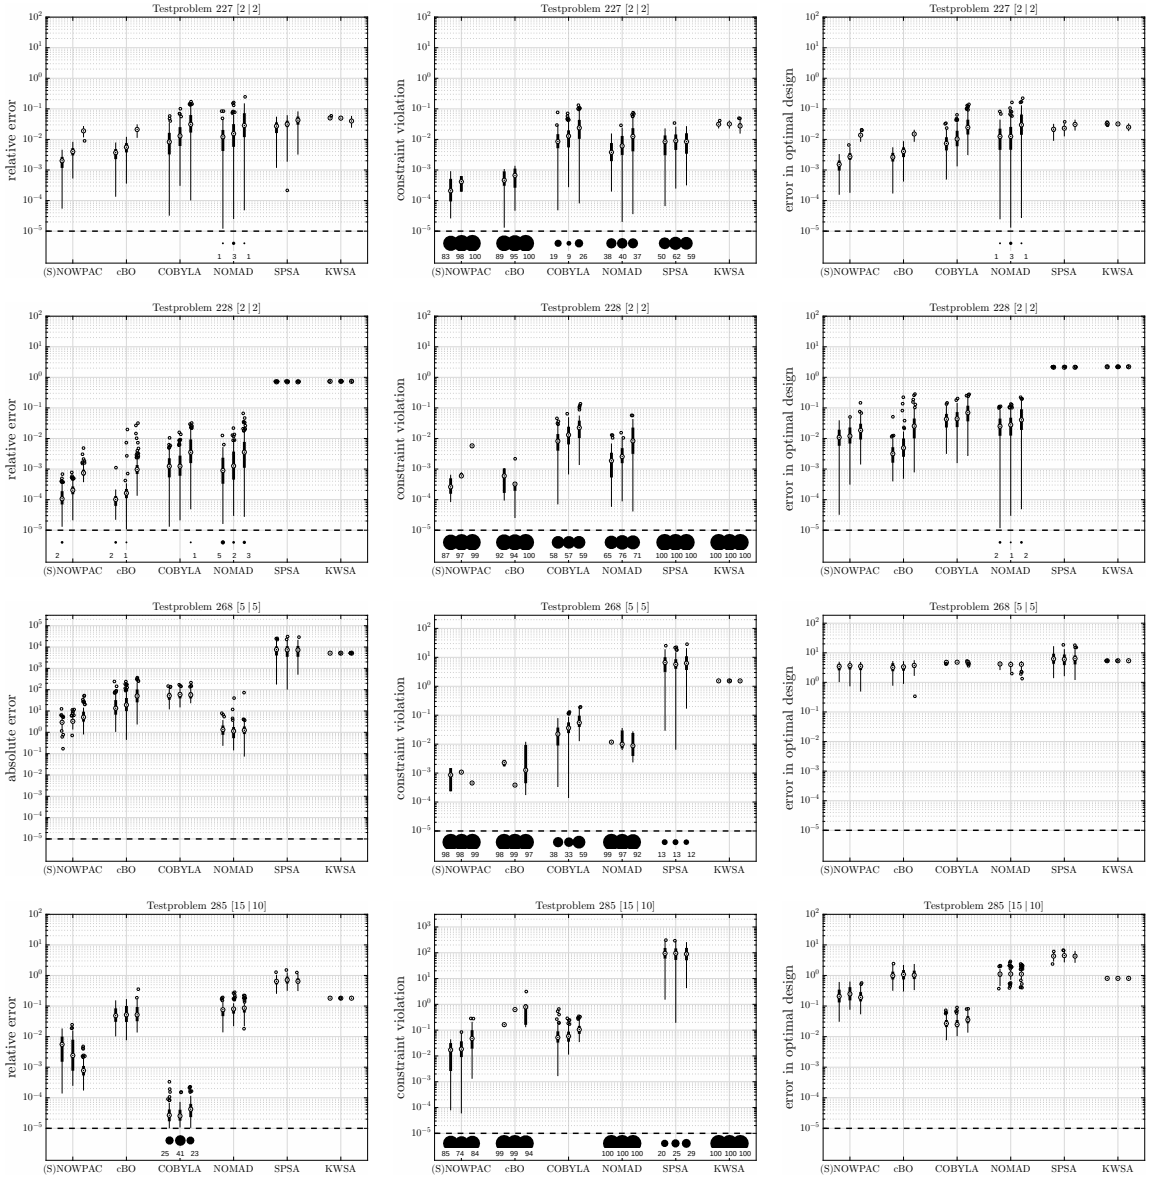

Figure 8: Box plots of the errors in the approximated optimal objective values (left plots), the constraint violations (middle plots) and the l_2 distance to the exact optimal solution (right plots) of 100 repeated optimization runs for the Schittkowski test problems number 227, 228, 268, and 285 for (22). The plots show results of the exact objective function and constraints evaluated at the approximated optimal design computed by (S)NOWPAC, cBO, COBYLA, NOMAD, SPSA and KWSA. Thereby all errors or constraint violations below 10^{-5} are stated separately below the 10^{-5} threshold and the box plots only contain data above this threshold.

or significantly better results than all the other optimization procedures.

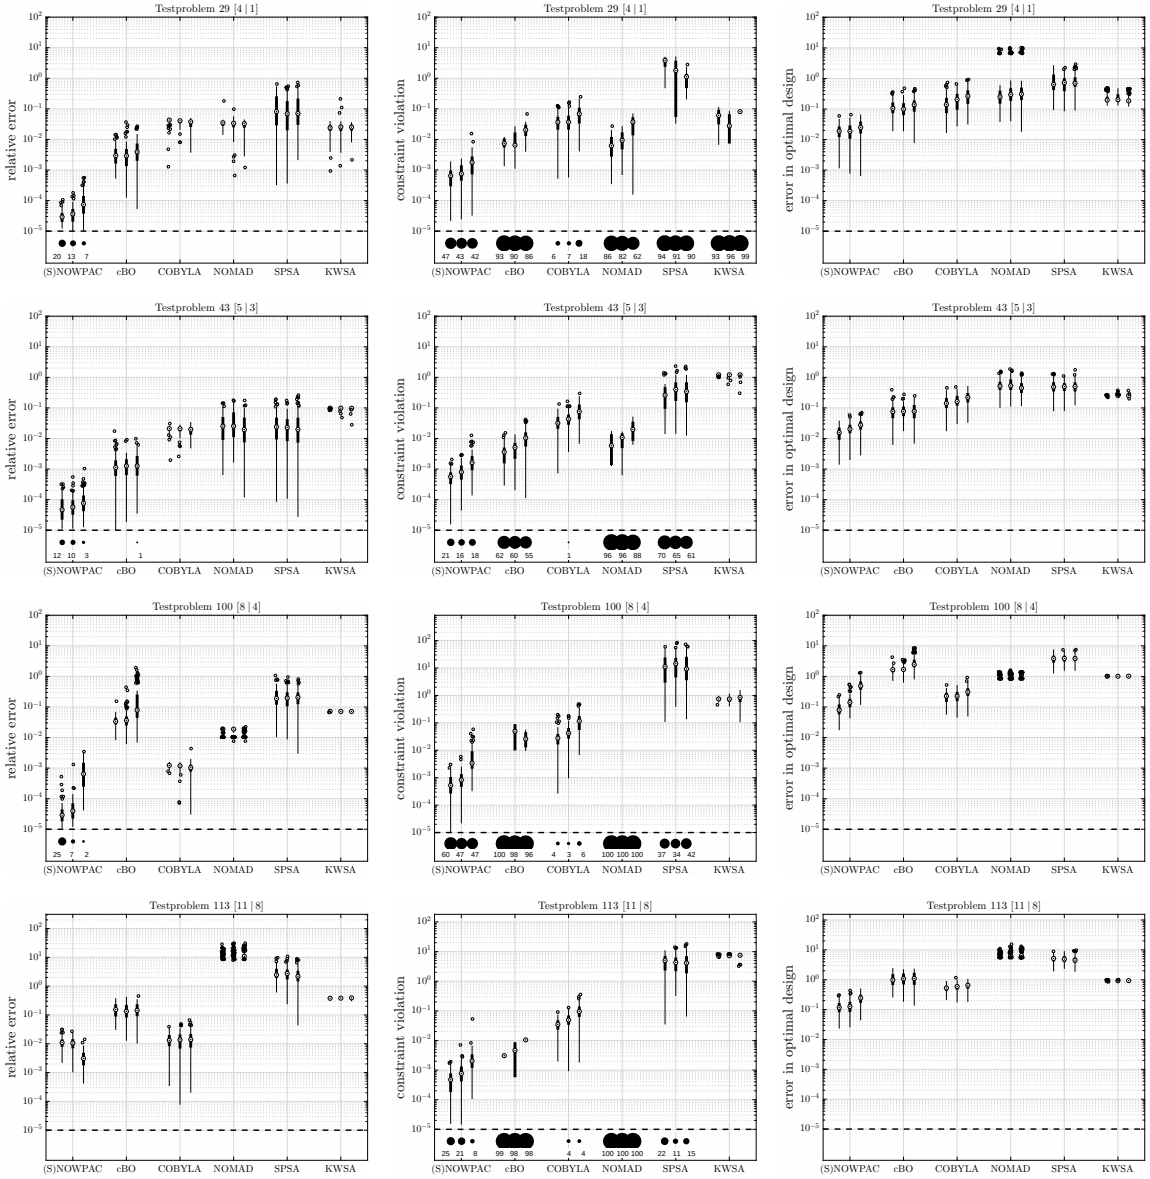

Figure 9: Box plots of the errors in the approximated optimal objective values (left plots), the constraint violations (middle plots) and the l_2 distance to the exact optimal solution (right plots) of 100 repeated optimization runs for the Schittkowski test problems number 29, 43, 100, and 113 for (23). The plots show results of the exact objective function and constraints evaluated at the approximated optimal design computed by (S)NOWPAC, cBO, COBYLA, NOMAD, SPSA and KWSA. Thereby all errors or constraint violations below 10^{-5} are stated separately below the 10^{-5} threshold and the box plots only contain data above this threshold.

5 Conclusions

We proposed a new stochastic optimization framework based on the derivative-free trust region framework NOWPAC. The resulting optimization procedure is capable of handling noisy black box evaluations of the objective function and the constraints.

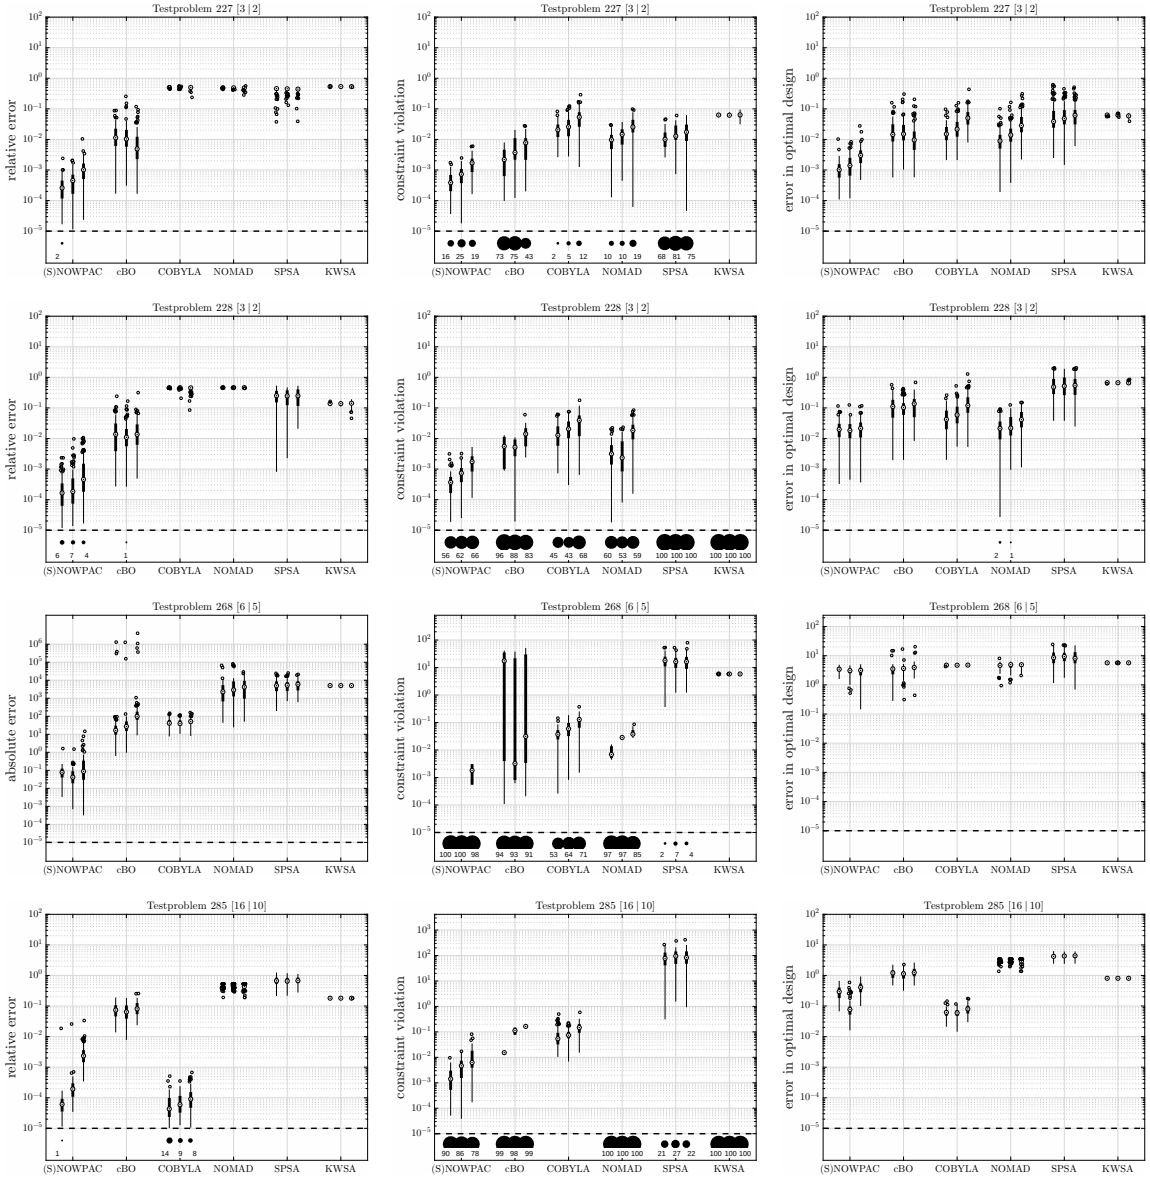

Figure 10: Box plots of the errors in the approximated optimal objective values (left plots), the constraint violations (middle plots) and the l_2 distance to the exact optimal solution (right plots) of 100 repeated optimization runs for the Schittkowski test problems number 227, 228, 268, and 285 for (23). The plots show results of the exact objective function and constraints evaluated at the approximated optimal design computed by (S)NOWPAC, cBO, COBYLA, NOMAD, SPSA and KWSA. Thereby all errors or constraint violations below 10^{-5} are stated separately below the 10^{-5} threshold and the box plots only contain data above this threshold.

Existing approaches for handling noisy constraints either rely on increasing accuracy of the black box evaluations or on Stochastic Approximation [79]. Increasing the accuracy of individual evaluations of the robustness measures may not be an efficient usage of computational effort as in local approaches individual black box evaluations are often discarded. We therefore introduced Gaussian process surrogates to reduce the noise in the black box

evaluations by re-using all available information. This is in contrast to Stochastic Approximation techniques [41] which only work with local gradient approximations, disregarding available information. Despite the rich convergence theory for Stochastic Approximation approaches, their practical application often strongly depends of the choice of technical parameters for step and stencil sizes as well as a penalty scheme for handling constraints. In our applications, without carefully tuning those parameters, Stochastic Approximation approaches performed suboptimal. Bayesian optimization techniques, in contrast make full use of all available data, resulting in computationally expensive optimization methods, in particular in higher dimensions. (S)NOWPAC combines the advantages of both worlds by utilizing fast local optimization with Gaussian process corrected black box evaluations. We showed in Section 4 that the overall performance of (S)NOWPAC is superior to existing optimization approaches approaches.

In our future work we will investigate convergence properties of our proposed stochastic derivative-free trust region framework towards a first order critical points.

Acknowledgments

This work was supported by BP under the BP-MIT Conversion Research Program. The authors wish to thank their collaborators at BP and the MIT Energy Initiative for valuable discussions and feedback on the application of (S)NOWPAC to energy conversion processes.

Bibliography

[1] C. Acerbi and D. Tasche. Expected shortfall: a natural coherent alternative to Value at Risk. *Economic Notes*, 31(2):379–388, July 2002.

[2] S. Ahmed and A. Shapiro. Solving chance-constrained stochastic programs via sampling and integer programming. In *Tutorials in Operations Research*. INFORMS, 2008.

[3] S. Alexander, T. F. Coleman, and L. Y. Minimizing CVaR and VaR for portfolio of derivatives. *Journal of Banking & Finance*, 30:583–605, 2006.

[4] P. Artzner, F. Delbaen, J.-M. Eber, and D. Heath. Thinking coherently. *Risk*, 10:68–71, 1997.

[5] P. Artzner, F. Delbaen, J.-M. Eber, and D. Heath. Coherent measures of risk. *Mathematical Finance*, 9(3):203–228, July 1999.

[6] C. Audet, A. L. Custodio, and J. E. Dennis Jr. Erratum: mesh addaptive direct search algorithms for constrained optimization. *SIAM Journal on Optimization*, 18(4):1501–1503, 2008.

[7] C. Audet and J. E. Dennis Jr. Mesh adaptive direct search algorithms for constrained optimization. *SIAM Journal on Optimization*, 17(1):188–217, 2006.

[8] C. Audet and J. E. Dennis Jr. A progressive barrier for derivative-free nonlinear programming. *SIAM Journal on Optimization*, 20(1):445–472, 2009.

[9] F. Augustin and Y. M. Marzouk. NOWPAC: A path-augmented constraint handling approach for nonlinear derivative-free optimization, arXiv:1403.1931v3, 2014.

[10] G. Bayraksan and D. P. Morton. Assessing solution quality in stochastic programs. *Mathematical Programming, Series B*, 108:495–514, 2006.

[11] A. Ben-Tal and A. Nemirovski. Robust convex optimization. *Mathematics of Operations Research*, 22:769–805, 1998.

[12] A. Ben-Tal and A. Nemirovski. Robust solutions of uncertain linear programs. *Operations Research Letters*, 25:1–13, 1999.

[13] D. Bertsimas, D. B. Brown, and C. Caramanis. Theory and applications of robust optimization. *SIAM Review*, 53(3):464–501, 2011.

[14] H. G. Beyer and B. Sendhoff. Robust optimization - a comprehensive survey. *Comput. Methods Appl. Mech. Engrg.*, 196:3190–3218, 2007.

[15] S. Bhatnagar, H. L. Prasad, and L. A. Prashanth. *Stochastic recursive algorithms for optimization*, volume 434 of *Lecture notes in control and information sciences*. Springer-Verlag London Heidelberg New York Dordrecht, 2013.

[16] D. M. Bortz and C. T. Kelley. *Computational methods for optimal design and control*, volume 24 of *Progress in Systems and Control Theory*, chapter The simplex gradient and noisy optimization problems, pages 77–90. de Gruyter, 1998.

[17] R. G. Carter. On the global convergence of trust region algorithms using inexact gradient information. *SIAM Journal of Numerical Analysis*, 28(1):251–265, February 1991.

[18] G. C. Cawley and N. L. C. Talbot. Preventing over-fitting during model selection via Bayesian regularisation of the hyper-parameters. *Journal of Machine Learning Research*, 8:841–861, 2007.

[19] K. H. Chang, L. J. Hong, and H. Wan. Stochasic trust-region response-surface method (STRONG) - a new response-surface framework for simulation optimization. *INF*, 25(2):230–243, 2013.

[20] T. D. Choi and C. T. Kelley. Superlinear convergence and implicit filtering. *SIAM Journal on Optimization*, 10(4):1149–1162, 2000.

[21] A. R. Conn, N. Gould, A. Sartenaer, and P. L. Toint. Global convergence of a class of trust region algorithms for optimization using inexact projections on convex constraints. *SIAM Journal on Optimization*, 3(1):164–221, February 1993.

[22] A. R. Conn, K. Scheinberg, and L. N. Vicente. Global convergence of general derivative-free trust-region algorithms to first- and second-order critical points. *SIAM Journal on Optimization*, 20(1):387–415, 2009.

[23] H. A. David and H. N. Nagaraja. *Order statistics*. John Wiley & Sons, Inc., Hoboken, New Jersey, 3rd edition, 2003.

[24] G. Di Pillo, S. Lucidi, and F. Rinaldi. A derivative-free algorithm for constrained global optimization based on exact penalty functions. *Journal of Optimization Theory and Applications*, Springer Science+Business Media New York(November), 2013.

[25] P. Frazier, W. Powell, and S. Dayanik. The knowledge-gradient policy for correlated nonorm beliefs. *INFORMS Journal on Computing*, 21(4):599–613, May 2009.

[26] J. R. Gardner, M. J. Kusner, Z. Xu, K. Q. Weinberger, and J. P. Cunningham. Bayesian optimization with inequality constraints. In *Proceedings of the 31st International Conference on Machine Learning*, 2014.

[27] R. B. Gramacy, G. A. Gray, S. Le Digabel, H. K. H. Lee, P. Ranjan, G. Wells, and S. M. Wild. Modeling an augmented Lagrangian for blablack constrained optimization. *Technometrics*, to appear, 2015.

[28] M. Heinkenschloss and L. N. Vicente. Analysis of inexact trust-region SQP algorithms. *SIAM Journal on Optimization*, 12(2):283–302, 2002.

[29] R. Henrion and A. Möller. A gradient formula for linear chance constraints under Gaussian distribution. *Mathematics of Operations Research*, 37(3):475–488, 2012.

[30] R. Hettich and K. O. Kortanek. Semi-infinite programming: theory, methods, and applications. *SIAM Review*, 35:380–429, 1993.

[31] W. Hock and K. Schittkowski. *Lecture Notes in Economics and Mathematical Systems*, chapter Test examples for nonlinear programming, no. 187. Springer, 1981.

[32] R. Hooke and T. A. Jeeves. "Direct search" solution of numerical and statistical problems. *Journal of the ACM*, 8(2):212–229, April 1961.

[33] D. R. Jones, M. Schonlau, and W. J. Welch. Efficient global optimization of expensive black-box functions. *Journal of Global Optimization*, 13:455–492, 1998.

[34] A. Kannan and S. M. Wild. Obtaining quadratic models of noisy functions. Technical Report ANL/MCS-P1975-1111, Argonne National Laboratory, 9700 South Cass Avenue Argonne, Illinois 60439, September 2012.

[35] C. T. Kelley. *Iterative methods for optimization*. SIAM, Society for Industrial and Applied Mathematics, Philadelphia, 1999.

[36] A. Kibzun and Y. Kan. *Stochastic programming problems: with probability and quantile functions*. John Wiley & Sons Ltd., 1996.

[37] A. Kibzun and S. Uryasev. Differentiability of probability function. *Stochastic Analysis and Applications*, 16(6):1101–1128, 1998.

[38] J. Kiefer and J. Wolfowitz. Stochastic estimation of the maximum of a regression function. *The Annals of Ma*, 23(3):462–466, 1952.

[39] S. Kim, R. Pasupathy, and S. G. Henderson. A guide to sample-average approximation. http://people.orie.cornell.edu/shane/pubs/SAAGuide.pdf.

[40] P. Krokhomal, M. Zabarankin, and S. Uryasev. Modeling and optimization of risk. *Surveys in Operations Research and Management Science*, 16:49–66, 2011.

[41] H. J. Kushner and G. G. Yin. *Stochastic approximation algorithms and applications*, volume 35 of *Applications of mathematics*. Springer Verlag New York, 1997.

[42] J. Larson and S. C. Billups. Stochastic derivative-free optimization using a trust region framework. *Computational Optimization and Applications*, 64(3):619–645, February 2016.

[43] J. Li, J. Li, and D. Xiu. An efficient surrogate-based method for computing rare failure probability. *Journal of Computational Physics*, 230:8683–8697, 2011.

[44] J. Li and D. Xiu. Evaluation of failure probability via surrogate models. *Journal of Computational Physics*, 229:8966–8980, 2010.

[45] J. Li and D. Xiu. Computation of failure probability subject to epistemic uncertainty. *SIAM Journal on Scientific Computing*, 34(6):A2946–A2964, 2012.

[46] G. Liuzzi, S. Lucidi, and M. Sciandrone. A derivative-free algorithm for linearly constrained finite minimax problems. *SIAM Journal on Optimization*, 16:1054–1075, 2006.

[47] G. Liuzzi, S. Lucidi, and M. Sciandrone. Sequential penalty derivative-free methods for nonlinear constrained optimization. *SIAM Journal on Optimization*, 20(5):2614–2635, 2010.

[48] A. March and K. Willcox. Constrained multifidelity optimization using model calibration. *Structural and Multidisciplinary Optimization*, 46:93–109, 2012.

[49] J. Mockus. On Bayesian methods for seeking the extremum. In *Optimization Techniques IFIP Technical Conference Novosibirsk*, volume 27 of *Lecture Notes in Computer Science*, pages 400–404. Springer-Verlag Berlin, 1974.

[50] J. Mockus. *Bayesian approach to global optimization: theory and applications*, volume 37 of *Mathematics and Its Applications*. Kluwer Academic Publisher Dordrecht, 1989.

[51] J. J. Moré and S. M. Wild. Benchmarking derivative-free optimization algorithms. *SIAM Journal on Optimization*, 20(1):172–191, 2009.

[52] J. A. Nelder and R. Mead. A simplex method for function minimization. *The Computer Journal*, 7(4):308–313, 1965.

[53] G. C. Pflug. *Optimization of stochastic models: the interface between simulation and optimization*. Kluwer Academic Publisher Boston, 1996.

[54] M. J. D. Powell. *Advances in Optimization and Numerical Analysis*, chapter A direct search optimization method that models the objective and constraint functions by linear interpolation, pages 51–67. Kluwer Academic, Dordrecht, 1994.

[55] M. J. D. Powell. Direct search algorithms for optimization calculations. *Acta Numerica*, 7:287–336, January 1998.

[56] A. Prékopa. On probabilistic constrainted programming. In *Proceedings of the Princeton Symposium on Mathematical Programming*. Princeton University Press, Princeton, NJ, 1970.

[57] J. Quinonero-Candela and C. E. Rasmussen. A unifying view of sparse approximate Gaussian process regression. *Journal of Machine Learning Research*, 6:1939–1959, 2005.

[58] R. Rackwitz. Reliability analysis - a review and some perspectives. *Structural Safety*, 23(4):365–395, October 2001.

[59] C. E. Rasmussen and K. I. Williams. *Gaussian processes for machine learning*. MIT Press, 2006.

[60] R. Reemtsen and J.-J. Rückmann, editors. *Semi-infinite programming*, volume 25 of *Nonconvex optimization and its applications*. Springer Science+Business Media Dordrecht, 1998.

[61] R. G. Regis. Stochastic radial basis function algorithms for large-scale optimization involving expensive black-box objective and constraint functions. *Computers & Operations Research*, 38(5):837–853, 2011.

[62] R. G. Regis. Constrained optimization by radial basis function interpolation for high-dimensional expensive black-box problems with infeasible initial points. *Engineering Optimization*, 46(2):218–243, 2014.

[63] H. Robbins and S. Monro. A stochastic approximation method. *Annals of Mathematical Statistics*, 22(3):400–407, 1951.

[64] R. T. Rockafellar and S. Uryasev. Optimization of conditional value-at-risk. *Journal of Risk*, 2(3):21–41, 2000.

[65] R. T. Rockafellar and S. Uryasev. Conditional value-at-risk for general loss distributions. *Journal of Banking & Finance*, 26:1443–1471, 2002.

[66] R. T. Rockafellar, S. Uryasev, and M. Zabarankin. Deviation measures in risk analysis and optimization. Technical report, Research Report 2002-7, Risk Management and Financial Engineering Lab, Center for Applied Optimization, University of Florida, 2002.

[67] R. Y. Rubinstein and A. Shapiro. *Discrete event systems*. John Wiley & Sons Chichester New York, 1993.

[68] P. R. Sampaio and P. L. Toint. A derivative-free trust-funnel method for equality-constrained nonlinear optimization. *Compuational Optimization and Applications*, 61(1):25–49, 2015.

[69] K. Schittkowski. More test examples for nonlinear programming codes. In *Lecture Notes in Economics and Mathematical Systems*. Springer, 1987.

[70] K. Schittkowski. 306 test problems for nonlinear programming with optimal solutions - user's guide. Technical report, University of Bayreuth, Department of Computer Science, 2008.

[71] A. Shapiro, D. Dentcheva, and A. Ruszczynski. *Lectures on stochastic programming.* Society for Industrial and Applied Mathematics and the Mathematical Programming Society, 2009.

[72] S. Shashaani, H. Fatemeh, and P. Raghu. ASTRO-DF: a class of adaptive sampling trust-region algorithms for derivative-free simulation optimization. *Optimization online*, 2015.

[73] J. C. Spall. Multivariate stochastic approximation using a simultaneous perturbation gradient approximation. *IEEE Transactions on Automatic Control*, 37(3):332–341, March 1992.

[74] J. C. Spall. Implementation of the simultaneous perturbation algorithm for stochastic optimization. *IEEE Transactions on Aerospace and Electronic Systems*, 34(3):817–823, 1998.

[75] W. Spendley, G. R. Hext, and F. R. Himsworth. Sequential application of simplex design in optimisation and evolutionary operation. *Technometrics*, 4:441–461, 1962.

[76] G. Szegö. Measure of risk. *Journal of Banking & Finance*, 26:1253–1272, 2002.

[77] S. Uryasev. Derivatives of probability functions and some applications. *Annals of Operations Research*, 56:287–311, 1995.

[78] S. Uryasev. *Probabilistic Constrained Optimization: Methodology and Applications*, chapter Introduction to the theory of probabilistic functions and percentiles, pages 1–25. Kluwer Academic Publishers, 2000.

[79] I.-J. Wang and J. C. Spall. Stochastic optimization with inequality constraints using simultaneous perturbation and penalty functions. In *Proceedings of the 42nd IEEE, Conference on decision and control*, December 2003.

[80] Y. Zhang. A general robust-optimization formulation for nonlinear programming. *Journal of Optimization Theory and Applications.*

[81] R. Zieliński. Optimal quantile estimators; small sample approach. Technical report, IMPAN, preprint 653, November 2004.

[82] R. Zieliński. Optimal nonparametric quantile estimators. towards a general theory. a survey. *Communications in Statistics - Theory and Methods*, 38:980–992, 2009.